

2

AD-A228 815

DTIC FILE COPY

Linguistic and Pragmatic Constraints  
on Utterance Interpretation

DTIC  
FILED  
NOV 14 1990  
S D

Elizabeth Ann Hinkelman

Technical Report 288  
May 1990

DISTRIBUTION STATEMENT A  
Approved for public release;  
Distribution Unlimited

UNIVERSITY OF  
ROCHESTER  
COMPUTER SCIENCE

90 11 9 027

# Linguistic and Pragmatic Constraints on Utterance Interpretation

by

Elizabeth Ann Hinkelman

Submitted in Partial Fulfillment  
of the  
Requirements for the Degree

Doctor of Philosophy

Supervised by James F. Allen

Department of Computer Science  
College of Arts and Sciences

JR 65 TR # 1

Dist. "A" per telecon Dr. Andre van  
Tilborg. Office of Naval Research/  
code 1133.

VHG

11/13/90

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By <i>per call</i>	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

University of Rochester

Rochester, New York

1989

## Curriculum Vitae

Elizabeth Ann Hinkelman was born in Rochester, New York on February 21st, 1963. She enjoyed geometrical toys in a Montessori-method kindergarten, and books thereafter. She learned Latin between the math and science at Nazareth Academy, a girls' high school, generating bad computer poetry at an Explorer post, on the side. She entered Colgate University in 1980, and was just trying to decide between the majors of neuroscience and philosophy when the ACL met in Toronto in '82. Her views of language became warped by exposure to stack-based models of discourse. She spent 1983 abroad in Switzerland, learning group theory on the Rhine at the University of Basel and microcomputer programming from the Master at the Eidgenossische Technische Hochschule, Zurich. It was there that she became permanently addicted to the academic lifestyle.

She graduated from Colgate in 1984 with various honors, a major in math and computer science, and several niggling questions left over from IJCAI '83. She began studying Artificial Intelligence at the University of Rochester that summer. Two years later her Master's Degree in Computer Science was awarded. She spent the summer of 1987 at BBN Labs, exploring the recognition of discourse plans. This dissertation was defended in September 1989.

## Acknowledgements

My advisor James Allen was a pervasive influence on this work. He taught me most of what I know about research and about natural language processing. Committee members Mike Tanenhaus, Greg Carlson, Len Schubert and Jerry Feldman made very interesting and constructive comments. Several additional people read drafts: Chung-Hee Hwang, Dave Traum, Steve Feist, and Massimo Poesio. Thank you; it's much better because of you.

The Computer Science Department at the University of Rochester provided a very congenial environment for this work, especially the administrative and technical staff. My parents hosted parties and endured excuses of busyness. And Rick believed when I didn't. Thank you all.

This work was supported in part by NSF research grants DCR-8502481 and IST-8504726, the US Army Engineering Topographic Laboratories research contract DACA76-85-C-0001, and ONR N00014-82-K-0193.

## Abstract

To model how people understand language, it becomes necessary to understand not only grammar and logic, but also how people use language to affect their environment. This area of study is known as natural language pragmatics. Speech acts, for instance, are the offers, promises, announcements, and so on that people make by talking. The same expression may be different acts in different contexts, and yet not every expression performs every act. We want to understand how people are able to recognize each other's intentions and implications in saying something.

Previous plan-based theories of speech act interpretation do not account for the conventional aspect of speech acts. They can, however, be made sensitive to both linguistic and propositional information. This document presents a method of speech act interpretation which uses patterns of linguistic features (e.g. mood, verb form, sentence adverbials, thematic roles) to identify a range of speech act interpretations for the utterance. These are then filtered and elaborated by inferences about agents' goals and plans.

In many cases the plan reasoning consists of short, local inference chains (that are in fact conversational implicatures), and extended reasoning is necessary only for the most difficult cases. The method is able to accommodate a wide range of cases, from those which seem very idiomatic to those which must be analyzed using knowledge about the world and human behavior. It explains how "Can you pass the salt?" can be a request while "Are you able to pass the salt?" is not.

## Table of Contents

1 Natural Language Pragmatics .....	1
1.1 Speech Acts .....	1
1.2 Conversational Implicature .....	3
1.3 Reasoning About Plans .....	8
1.4 Previous Work .....	11
1.5 Overview of a New Approach .....	16
1.6 Overview of the Thesis .....	18
2 Linguistic Constraints I: Fundamentals .....	20
2.1 Representation of Linguistic Structures .....	22
2.2 Evidence for Interpretations .....	24
2.3 Applying the Rules .....	32
2.4 Combining Partial Descriptions .....	36
2.5 Discussion .....	38
2.5.1 Another Example .....	42
3 Linguistic Constraints II: Case Studies and Limits .....	46
3.1 Sentence Type .....	46
3.2 <i>Hereby</i> .....	50
3.3 <i>Please</i> .....	53
3.3.1 A Little Etymology .....	53
3.3.2 Uses of <i>Please</i> .....	55
3.3.3 The Data .....	58
3.3.4 Taxonomy of Speech Acts .....	63
3.4 Syntactic Complications .....	65
3.5 The Limits of Conventionality .....	66
4 Plan Reasoning I .....	69
4.1 Role of Plan Reasoning .....	69
4.1.1 Introduction .....	69
4.1.2 Plans and Discourse .....	70
4.1.3 Plan Reasoning with Speech Acts .....	71
4.1.4 Discussion .....	77
4.2 Related Work .....	81
4.2.1 Perrault .....	81
4.2.2 Cohen and Levesque .....	84
4.2.3 Kautz .....	85
4.3 Short Inferential Distance .....	86
4.3.1 Plan-based Conversational Implicature .....	87
4.3.2 Implicature and Speech Acts .....	91
5 Plan Reasoning II .....	94
5.1 Knowledge Representation .....	94

5.1.1 The Logic .....	95
5.1.2 Notation for Actions .....	97
5.1.3 Plan-Based Implicatures .....	99
5.2 A Speech Act Taxonomy .....	101
5.3 Some Simple Examples .....	111
5.3.1 Speech Act Preconditions .....	111
5.3.2 Negations of Effects .....	114
5.3.3 Intended Effects .....	115
5.3.4 Constraints .....	116
6 Two Constraints Integrated .....	118
6.1 Interaction of the Constraints .....	119
6.2 An Extended Example .....	121
6.2.1 "Can you speak Spanish?" -- Context One .....	122
6.2.2 "Can you speak Spanish?" -- Context Two .....	125
6.2.3 "Can you speak Spanish, Please?" -- Context One .....	127
6.2.4 "Can you speak Spanish, Please?" -- Context Two .....	129
6.2.5 Comparison .....	129
6.3 Extended Reasoning .....	131
6.3.1 The Interface .....	131
6.3.2 An Example of Extended Reasoning .....	136
7 Implementation .....	140
7.1 The Rochester Discourse System .....	140
7.2 The Linguistic Module .....	142
7.3 The Implicature Module .....	147
7.4 Limitations .....	150
8 Conclusion .....	151

## 1. Natural Language Pragmatics

Whether people use language grammatically or ungrammatically, accurately or inaccurately, they are using it to realize their own goals or intentions. In philosophy and linguistics, this aspect of language is referred to as pragmatics.

Consider a brief encounter between two strangers, from [Grice 75]:

A is standing by an obviously immobilized car and is approached by B;  
the following exchange takes place:

- (1) A: I am out of petrol.  
B: There is a garage round the corner.

B communicates that the garage is open and has gas to  
sell, and so on.

These implications arise because we believe that B is trying to make a helpful suggestion, not merely spouting random propositions. We would like to know precisely how hearers recognize each other's intentions. We want to know in what sense A's utterance is a request for help, and B's is a suggestion. We want to know how the various implications are made. We must show how an agent's use of language for specific goals is related to traditional subjects of language study like syntax and semantics.

### 1.1. Speech Acts

An utterance is a small unit of linguistic output, a sentence or fragment, by a particular person in a particular situation. The notion that utterances are actions rather than merely descriptions is due to Austin [Austin 62]. Sentences like

- (2) a: I hereby dub thee Knight.



- b: I promise to be home by midnight.
- c: I name this ship the *Queen Elizabeth*.
- d: I affirm that this information is true, to the best of my knowledge.

when uttered sincerely and with authority, constitute a social event. Dubbing is felicitously performed by persons of a certain social rank in a certain culture, with ceremony and sword-waving, when they wish to bestow the rank of knight on an inferior. A promise like (b) is a domestic event which might occur between a teenager and parent, when one is planning to go out for the evening. The reader can imagine a context for (c) and (d). These syntactically rigid sentences, uttered in context, are referred to as explicit performative utterances. They are prototypical of linguistic actions, which Austin called *speech acts*. They may express attitudes, as greetings do, commitment, like promises, information, like assertions, judgement, as in sentencing, or attempts to get someone to do something, like requests and commands. Many such actions can be carried out in nonlinguistic ways as well.

The problem of so-called indirect speech acts [Searle 75] concerns sentences like

- (3) a: Can you pass the salt?
- b: You're standing on my toe.
- c: Has anyone offered you a ride to the airport?

At first glance, (a) is a yes/no question, (b) is a statement, and (c) is another yes/no question. Yet in some common contexts, (a) is a request, (b) is a request, and (c) is an offer. There is no simple mapping between sentence form and speech act type. One also has the sense that, unlike idioms, these sentences often seem to have both interpretations at the same time. We would like to explain how the

speech acts can be identified, and how they are related to the literal meaning. We must be careful not to underestimate the richness of the contexts in which the utterances occur, despite their familiarity.

Searle's proposal was to relate the propositional content of the sentence to the intended speech act via the appropriateness or *felicity conditions* of that type of action. For instance, it is only felicitous to request actions which the hearer can perform: therefore since (a) asks if the hearer can pass the salt, it may be a request to pass the salt. [Perrault 80] developed a computational version of these ideas, based on Artificial Intelligence models of reasoning about actions.

However, several kinds of information complicate the recognition process. Certain words tend to be associated with certain speech act types, and sentence mood and other syntactic features play a role too. Literal meaning, lexical and syntactic choices, agents' beliefs, the immediate situation, and general knowledge about human behavior all clarify what the speaker's intentions are. The present work shows how these factors can be integrated into a model of speech act interpretation which handles the full range of speech acts in a clean way.

## 1.2. Conversational Implicature

Grice's problem of conversational implicature, illustrated by the gas station example, is closely related and indeed, overlapping. In order to know what final conclusions to draw from an utterance, we need to know initially what action is being done. Both problems require a logic for modelling human action, allowing

the hearer to model the speaker's reasoning. Conversational implicatures are by definition *cancellable*: they are neither part of the sentence's truth conditions nor its entailments with respect to some set of universal rules, and so may be denied without contradicting the sentence. For instance, an alternative response by B is unhelpful but not contradictory:

- (4) A: I'm out of petrol.  
B: There's a garage round the corner, but it is closed.

They are *nondetachable*: the same implicatures are associated with any paraphrase of the utterance.

- (5) A: I'm out of petrol.  
B: If you go around the corner, you'll find a garage.

In practice the paraphrase test sometimes fails, but the idea is that implicatures are reasoned from the propositional content of the utterance and not from lexical connotations. Conversational implicatures may also be open or indeterminate, if the context suggests more than one possibility, or if it is unclear exactly what is being suggested. Consider this conversation in a car on the interstate:

- (6) Pat: Wanna get off at the next exit for dinner?  
Sandy: It's fairly early yet....

Is Sandy saying that it's too early to eat? That it would be good to go now and avoid the crowds? That there is plenty of time to eat now? Grice's point is that the indeterminacy here is part of the phenomenon, rather than a failure of the theory. The hearer may have several incompatible conclusions with no way to distinguish among them.

Hirschberg [Hirschberg 85] considered a class of examples which make use of some underlying set of values, such as the following:

- (7) Chris: Is there a department store nearby?  
Dana: K-Mart is probably open.

Dana implicates that there are other stores nearby, but that they may not be open. Hirschberg stressed the role of sets of values, and any orderings that apply to them, in her theory of *scalar implicature*. She showed that with a few general rules, it is possible to draw implicatures from utterances relying on such diverse partial orderings as the colors, possible baseball scores, modal verbs ordered by degree of possibility or degree of obligation, the steps in a process, and temperatures. Sentences like

- (8) She should be home by now.

in which *should* represents an intermediate certainty value, imply that lower values are true (she could be home now) and that higher values (she is definitely home now) are false or unknown. Likewise

- (9) It's not warm out.

normally implies that it isn't hot out, and that colder values (chilly, cold, freezing) are true or unknown. Of course temperature can be viewed in terms of warmth rather than coolness, the ordering running in the other direction:

- (10) It's finally warm out.

Here we know that it's not freezing or cold, and it may even be hot. Hirschberg's work is the first implicature model specific enough to be implemented as a computer program.

[Hinkelman 87] considered *plan-based* implicatures. The key to this computational model is to re-examine the gas station example, using a model of human action. Speech acts and domain acts are represented as plans, structured objects consisting of preconditions, steps, and effects. Each aspect of plan representation becomes a basis for certain inferences. In the gas station example, we all know that the ordinary way to get gas is to go to a gas station and pump it and pay for it. There are variations in this situation: A cannot drive there and will have to collect the gas with a gas can. But either way, A and B both know that for A to get gas, the gas station must be open, have gas, etc. These are preconditions of buying gas. Someone will have to pump the gas, and A will have to pay for it, presumably: these are steps of the plan. In the end, A will have the gas and be able to drive on. These are effects of buying gas, in this situation. B implicates that, as far as B knows, all of this is true. Otherwise the suggestion would be unhelpful. An argument for this approach was made in the philosophy literature by [McCafferty 86].

In the study of communication, great care must be taken to distinguish beliefs of the speaker, beliefs of the hearer, and shared beliefs. If A successfully informs B of some fact, one result is that B now not only believes the fact, but believes that the fact is a shared belief, and further, that A believes that the fact is a shared belief. Likewise, A now believes that B believes the fact, and that B believes it's

mutually believed, and so on. Such accounting is needed to explain communication failures and lying. Here is a misunderstanding from real life:

A student who was thinking about buying some candy from a vending machine went to the library desk.

- (11) Student: Can I have change? (proffers \$20.)  
 Librarian: Not for a twenty.

The student went and found a friend, who traded two tens, but was told on returning that in fact change was available only for the photocopy machines.

The librarian implicated that the student could have change for smaller bills, perhaps assuming that the student intended to make copies. The student inferred that he could have change for smaller bills, and that the librarian intended them both to believe this, but had no way of inferring the restriction to copying. They came to a mutual belief that his plan would work, but with different beliefs about what the plan was. Under both interpretations the literal content of the librarian's statement is true, but the exact implicatures are different because the plan is different. Had the librarian correctly recognized the student's plan, he would have been obliged to state that he was unable to change any denomination.

Given an utterance and context, we model how the utterance changes the hearer's belief state. Recognition of the speech act and recognition of the implicatures are tightly bound subproblems in this task. This thesis reinforces the claim that a theory of human action is important in understanding language phenomena, and extends its scope somewhat. Its main contribution along that line will be to show

how such a theory of action must interact with linguistic factors to provide broad-coverage speech act interpretation.

### 1.3. Reasoning about Plans

We must briefly introduce plan-based speech act interpretation here, to show that it fails to account for linguistic constraints on speech act interpretation.

Typical components of plan reasoning include a library of stored plans, some rules and algorithms to use in constructing plans (planning), rules and algorithms for inferring other agents' plans (plan recognition), and a knowledge base of the agents' beliefs about the world and each other (context and world knowledge). These components of an intelligent social agent provide the basis for pragmatic interpretation of utterances as well.

The representation of actions that will be discussed here is in the tradition of early work on planning, exemplified by the STRIPS system [Fikes 71, Nilsson 80, Sacerdoti 74, Sacerdoti 80]. Here actions are operators on a database. Their descriptions include a set of propositions which describe the conditions under which the action can succeed, called preconditions. We subdivide these into true preconditions, which the agent can plan to achieve, and constraints, which cannot be effected by the agent. (A historical note: constraints were invented for to solve the technical problem of ensuring that variables in different substeps kept the same bindings throughout.) There are also *add* and *delete* lists, propositions which result from an action performed when its preconditions hold.

We collapse these into the category of effects. An action can be semantically interpreted as an operator mapping the set of possible worlds described by its preconditions into the set described by its effects, consistent with its variable bindings. An action token has its agent and type parameters bound; an action type does not. In this document an action in a hierarchy or definition is always an action type, and in an example it is always an action token, although we will make no notational distinction.

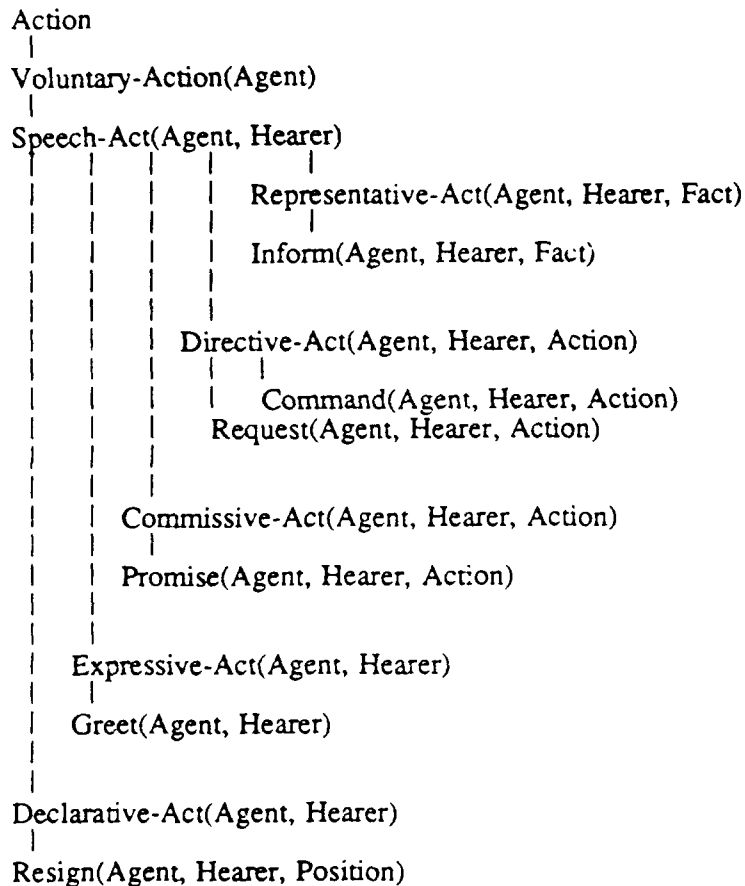
An example of such an elementary action type is `MOVE(Agent, Loc1, Loc2)`. Such an action might be defined with the precondition that Agent is in location Loc1, and constraint that Agent is a functioning animate being. The effect of this action is that Agent is in Loc2. The `MOVE` action type here is primitive or *basic* in the sense that it has no component actions. Nonbasic actions have a body which consists of other actions, which may have ordering constraints, and these must ultimately be decomposable into basic actions. We will refer to such actions as plans. Properly a plan token includes the initial and final conditions as well as the sequence of actions that makes this transformation, but in practice we will use the terms plan and action interchangeably. We also define an abstraction hierarchy on the action types. The abstraction relation states that if type T abstracts type T', any action A satisfying type T' also satisfies type T.

The figure below sketches an abstraction hierarchy for speech acts, denoted only by their types and parameters. The class of speech acts is a subtype of voluntary actions, and it subdivides into five main categories taken from [Searle 79]. For



each category we give just one example subtype, although there are many. Representative acts are those in which the speaker asserts some description the world's state, regardless of the degree of belief, or of the accuracy of the description. Informing, hypothesizing, and boasting fall into this category. Directive acts are attempts to get the hearer to do something; requests and commands are the paradigm examples. Commissive acts are those in which the speaker is bound to bring about a state of the world, and promises are prototypical commissives. (We may occasionally refer to the entire class by mentioning a prototypical example.) Expressive acts, such as condolences, are nominally expressions of attitude about some state of affairs, and not in general attempts to achieve something or describe the world. Searle contrasts them with declarative acts, which comprise the institutional explicit performative acts like "You're fired!". Such acts do generally create the state of affairs that they mention. In this abstraction hierarchy, if a Greet act is successfully performed, a Speech-Act has occurred with all of its preconditions and effects. Plan reasoning systems differ in the exact relations represented, but in general all this information is available in some form.

In STRIPS-like systems, planning is a process of chaining together actions by matching preconditions with effects. It can be viewed as search through the space of possible action sequences. In plan recognition, observed actions are used to identify the plans they may be a part of, and the goals to be met by those plans. Plan execution traces through a predefined plan, in chronological order.



Understanding speech acts and implicatures may require utilizing any of these sorts of reasoning, which are independently needed by intelligent agents.

#### 1.4. Previous Work

Previous work on speech act interpretation falls roughly into three approaches, each with characteristic weaknesses: the idiom approach, the plan based approach, and

the descriptive approach.

The idiom approach is motivated by pat phrases like

- (12) a: Can you please X? (request, literally a yes/no question)
- b: Would you kindly X? (request, literally a yes/no question)
- c: I'd like X. (request, literally an inform of hypothetical attraction)
- d: May I X? (request, literally a yes/no question)
- e: How about X? (suggestion, literally a question.)

The system could look for these particular strings, and build the corresponding speech act using the complement as a parameter value. If this simple method were effective, speech act interpretation would be uninteresting. Something similar was proposed in [Lehnert 78].

- (13) a: Do you know X?
- b: Tell me X.

Lehnert takes parsed sentences of the form (a) and substitutes semantic representations to get (b), then processes the new sentence further. But such sentences are not true idioms, because the literal meaning also plays a role in many contexts. One can respond to the literal *and* nonliteral acts: "Yes, it's the 9th." The idiom approaches are too inflexible to choose the literal reading or to accommodate ambiguity. They lack a theory connecting the nonliteral and literal readings.

Another problem is that some classic examples are not even pat phrases:

- (14) a: It's cold in here.
- b: Do you have a watch on?

In context, (a) may be a request to close the window. Sentence (b) may be asking what time it is or requesting to borrow the watch. Handling sentences like these

requires extensive ability to reason about plans.

The plan based approach [Allen 83, McCafferty 86, Perrault 80, Sidner 81] [Brown 80] presumes a mechanism modelling human problem solving abilities, including reasoning about other agents and inferring their intentions. The system has a model of the current situation and the ability to choose a course of action. It can relate uttered propositions to the current situation: being cold in here is a bad state, and so you probably want me to do something about it; the obvious solution is for me to close the window, so, I understand, you mean for me to close the window. The plan based approach provides a tidy, independently motivated theory for speech act interpretation.

It does not use language-specific information, however. Consider

- (15) a: Can you speak Spanish?
- b: Can you speak Spanish, please?

The first sentence is a yes/no question in typical circumstances, but simply appending the word "please" forces the interpretation to a request. This is not peculiar to "please". The (a) sentence below may be a question or a request, yet paraphrases (b)-(d) are not requests.

- (16) a: Can you open the door?
- b: Are you able to open the door?
- c: Are you capable of opening the door?
- d: I hereby ask you to tell me if you can open the door.

In the following sets of sentences, only the first is a possible request; the paraphrases are not, unless sarcastic ones.

- (17) a: Would it be possible for you to open the door?  
b: Is it possible for you to open the door?
- (18) a: Why don't you open the door?  
b: How come you don't open the door?  
c: What's the reason that you don't open the door?
- (19) a: Do you mind opening the door?  
b: Are you opposed to opening the door?

(17a is most commonly a suggestion, but it can be a request.) Further, different languages realize speech acts in different ways. These examples, from [Sadock 74], are taken from Swedish, Hebrew, and Greenlandic, respectively.

They are followed by their literal translations.

- (20) a: Tank om Ni skulla opna doren.  
b: Think whether you should open the door.
- (21) a: ata muxan liftoax et hadelet?  
b: Are you ready to open the door?
- (22) a: matumik angmarniarit  
b: May you try to open the door.

Here is a different example, where (a) is translated from Hebrew:

- (23) a: You want to cook dinner.  
b: You wanna toss your coats in there?

The declarative sentence (a) can be a request, idiomatic to Hebrew, while the nearest American expression is interrogative (b). Neither is a request in British English.

- (24) a: Can you hand me that book?  
b: Muzete mi podat tu Knizku?

According to Searle, our (a) is very odd as a request in Czech (b). Specific social acts often have very rigid forms, e.g., greetings (or see [Horn 84]).

- (25) a: Gruess Gott!  
b: Greet God!

This commonplace Bavarian greeting is not idiomatic when translated literally. And speech acts vary with idiolect, too. Otherwise very cooperative persons may simply expect genuine requests to be stated explicitly. They simply do not recognize indirect requests. The plan based approach has nothing to say about these differences. Neither does it explain the psycholinguistic [Gibbs 84] finding that people access idiomatic interpretations in context more quickly than literal ones. Psycholinguistically plausible models cannot derive idiomatic meanings from literal meanings.

Descriptive approaches cover large amounts of data. [Brown 80] recognized the diversity of speech act phenomena and produced the first computational model with wide coverage. A representative rule from her system is *Equi-Ask*. It states that asking whether a particular speech act has been performed is a way of actually performing it.

- (26) a: Has anyone asked you to take out the trash? (request)
- b: Has anyone offered you a ride to the airport? (offer)
- c: Has anyone suggested Gerard Manley Hopkins? (suggestion)

Although it relied on a representation for actions, this proposal made few theoretical contributions. It also did not handle the language-specific cases well.

[Gordon 75] discuss sentences which are *sincerity conditions* of the speech act they perform. Sincerity conditions are similar to preconditions but are stated very generally: the speaker must believe what is said, the speaker can only request feasible actions, and so on. Gordon and Lakoff do not provide any criteria or motivation for what makes a good sincerity condition. Lacking a theory of human

action, they are also unable to explain utterances that rely on aspects of the requested or domain action, as in the Lehnert example. There one asks a question by asking literally whether the hearer knows the answer. A plan-based approach would argue that knowing the answer is a precondition for stating it, and this logical connection enables identification of the real question. The "Equi-Ask" construct is another example that does not fit readily into their framework. Their discussion of transderivational rules allowing interaction of syntax and pragmatics is suggestive but insufficiently explained.

### 1.5. Overview of a New Approach

We augment the plan-based approach with a linguistic component. The linguistic component consists of rules associating linguistic features with partial speech act descriptions. The rules express linguistic conventions that are often motivated by planning theory. They allow for an element of arbitrariness in just which forms are idiomatic to a language, and just which words and features mark it. They also allow for an interpretation process paralleling syntactic and semantic interpretation, with the same provisions for merging of partial interpretations and postponement of ambiguity resolution. The plan reasoning mechanism has none of these capabilities, nor have previous approaches. We will refer to the process of unifying several partial interpretations (versus a full interpretation from a single, more complex rule) as incremental, since each rule constrains the interpretation independently of the others.

Once the utterances have been interpreted by our conventional rules to produce a set of candidate conventional interpretations, these interpretations are filtered by the plan reasoner. Plan reasoning processes unconventional forms in the same spirit as earlier plan-based models, handling the same range of cooperative behavior from more refined input. We use a restricted version of plan reasoning for the ordinary cases, one which yields plan-based conversational implicatures as a bonus.

Consider what happens to an utterance as it passes through the system. Let us suppose, for the sake of concreteness, that a person named Suzanne is at the Spanish consulate, doing her paperwork for a Fulbright scholarship year in Spain. The representative, one Mrs. de Prado, asks

(27) Can you speak Spanish?

The system performs lexical and syntactic analysis of the sentence, and semantic interpretation. The linguistic component of speech act interpretation then generates a range of possible interpretations. It does this by attempting to match patterns of linguistic features against the analyzed sentence, each of which constrains the possible interpretations. Subject-auxiliary inversion suggests that this could be a yes-no question. The modal auxiliary with the hearer as subject suggests a request. Other patterns yield further constraints and very general interpretations. The sets of partial interpretations are combined incrementally to yield a request, yes-no question, and one more general interpretation.



Reasoning about plans is then used to filter the possible interpretations and constrain them further. Although in general it is possible to use the full power of methods like that of Allen and Perrault, we suggest that a more limited version is appropriate to ordinary cases. Our more limited version resembles a single breadth-first ply of the kind of axioms they used, and we show that the results are a class of conversational implicatures. The system computes this set of implicatures for each of the interpretations given above, and checks them for consistency with the hearer's other beliefs. Inconsistent interpretations are rejected and consistent ones favored. Remaining ambiguity may be resolved if necessary using extended reasoning or by generating further questions.

One implicature of the yes-no question, for instance, is that the speaker does not know the answer. (Didactic and rhetorical questions would have a different speech act type.) If Suzanne believes that Mrs. de Prado knows she speaks Spanish, she will eliminate the possibility of a sincere yes-no question. If she believes Mrs. de Prado does not know, she may be inclined to accept this possibility and its implicatures. If Suzanne is unsure, she may plan to address both of these possibilities in some way, or seek to disambiguate.

## 1.6. Overview of the Thesis

This thesis makes several contributions to the area of natural language pragmatics. It argues that both linguistic information and information about actions are necessary for a full account of speech acts. It presents a method of generating

speech act interpretations that makes full use of the linguistic description of the utterance. The method uses incremental rules and integrates readily with reasoning about plans. Reasoning about plans is explored also, yielding a place for conversational implicature in the architecture of natural language processing and defining a new and useful class of conversational implicatures. It is also shown what roles more extended plan reasoning may have in natural language processing. The system overall can be viewed as imposing several sets of constraints on utterance interpretation, with the input feeding up through them.

The structure of the thesis is this: Chapter Two explains the linguistic constraints on speech act interpretation, and the incremental pattern-matching method that embodies them. Chapter Three contains further examples of the method and discussion of the more complicated and limiting cases. Chapter Four introduces the plan reasoning aspect of utterance understanding, providing a preliminary specification of its functionality. Chapter Five construes plan reasoning as conversational implicature, making it sensitive to linguistic constraints and showing its general usefulness in utterance understanding. Chapter Six explains the interaction of the linguistic and pragmatic constraints on utterance interpretation, and the role of ambiguity. Chapter Seven describes the implementation, and Chapter Eight concludes the work.

## 2. Linguistic Constraints I: Fundamentals

We have seen that the speech act type is not an immediate function of sentence semantic content, nor is it simply a function of more extended inference. Although we can generally devise a post-hoc logical account of a particular utterance, there are numerous interactions with linguistic processing that must be accounted for. A final theory of speech acts must explain how people make use of lexical, syntactic, and semantic resources in expressing and recognizing intentions. Such a theory must show how this process is sensitive to paraphrase, to idiolect, and to the idiomatic aspects of the language being used.

Psycholinguistics has also suggested that literal meaning is not used in interpretation of indirect speech acts. Gibbs [Gibbs 84] argues that the distinction between literal and metaphorical meanings has no psychological equivalent. In particular, in context "indirect" speech acts are identified too quickly to involve the computation of literal meaning first. Neither can the literal meaning be a simultaneous calculation, since it fails to prime subsequent tasks based on it. Gibbs [Gibbs 86] also found that although subjects preferred to generate "indirect" requests corresponding to perceived obstacles to the request, the surface form expressing a particular obstacle was relatively fixed. Some tendency to favor shorter forms was observed, but no final conclusions about the favored forms were possible. Thus, from a psychological standpoint the role of surface elements in speech act interpretation is detectable, while a literal meaning phenomenon is not. Thus, conventions of language use [Morgan 75] have a large psychological role.

There are several phenomena whereby the speech act seems to intrude into sentence syntax. We present them here solely as an argument for linguistic processing of speech acts. The most obvious example is an explicit performative utterance of the "I hereby promise...." variety, where the main verb of the sentence may be taken to indicate the speech act. However, [Davison 83] reports on cases where it would be useful to assume the presence of a performative verb in the deep structure, because other sentence elements appear to modify such an item although it fails to appear on the surface. Clauses of manner and reason can have this property.

- (28) a: I'm just going to the store, in case you call and I'm out.  
 b: Andrew isn't here, because he isn't feeling well.

In (a) the reason clause is a reason for *stating* the main clause, not a reason for going to the store. This is in contrast to (b), where the reason simply modifies the main clause contents. One is very tempted to propose that the go-clause is dominated by a verb of stating, but this leads to difficulties which we will discuss later. Many sentential adverbs such as *frankly*, *strictly*, *confidentially* also modify the *stating* rather than the contents of the utterance. Adverbial phrases can also have this property, and there are questions of quantifier scope that appear to interact with a speech act marker. A complete theory of speech acts should also explain these phenomena.

The linguistic component of our model is the subject of this chapter. It will consist of incremental, language-specific rules which provide evidence for a set of

partial speech act interpretations. Later, we use plan reasoning to constrain, supplement, and decide among this set.

## 2.1. Representation of Linguistic Structures

Our notation is based on that of [Allen 87]. It incorporates lexical, syntactic, and semantic information about sentences. Its essential form is a parenthesized list, consisting of a category name followed by any number of slot/filler pairs. The syntactic component of the representation has conventional syntactic categories like S or NP. Syntactic slots correspond to subconstituent roles like subject, or features. A filler may be a word, a feature value, or or another (category...) structure. If a feature value appears in a subconstituent slot, it restricts the final filler of that slot to be a unit having that feature value. Alternation is represented by a list of possible values in curly brackets. Thus, the syntactic fragment below has category S for sentence, a slot for sentence mood with the value yes/no question, and a subject subconstituent. The subject has category NP, a head slot containing the word "you", and a number slot which may be singular or plural.

```
(S MOOD YES-NO-Q
  SUBJ (NP HEAD you
        NUM {s p}))
```

We divide semantics into two parts. The first, logical form, is used to capture the linguistic generalities of verb subcategorization and noun phrase structure. It embodies the hypothesis that a small, finite set of thematic roles is sufficient to

explain the semantic phenomena of linguistics [Carlson 84, Jackendoff 72]. Semantic categories are much more specialized than syntactic ones, including types of actions, states, and objects. Semantic slots are tense, and thematic roles such as agent, object, instrument, and from-location. Semantic fillers are constants identifying semantic objects, or (category...) structures. Semantic structures reside in SEM slots. The second component of semantics is the representation language used by the knowledge base, which resembles frame-based languages and has an unrestricted set of roles that range from very general to very specific. It represents actions and states, incorporating selectional restrictions, identification of referents, and other phenomena involving world knowledge. Knowledge base classes may be very abstract or very specific classes of actions, states, and objects. They are less restricted than thematic roles both in their total number and in the number that any instance may have. Knowledge base slots are, as we said, more detailed roles. There is a certain amount of commonality and even common terminology between logical form and knowledge base slots. Knowledge base fillers are knowledge-base objects (referents), which may recursively have internal structure. Knowledge base structures appear in REF slots. It is important to our method that the components are all available to the pragmatic interpretation process, and so for simplicity of presentation we will allow logical form and knowledge representations to appear in slots on the syntactic structure. The mapping between SEM and REF structures is not an issue that we can address here (although it is a linguistic computation), nor is the problem of reference.

Our representation of the sentence "Can you speak Spanish?" is shown below.

```
(S MOOD YES-NO-Q
  VOICE ACT
  SUBJ (NP HEAD you
        SEM (HUMAN ID h1)
        REF Suzanne)
  AUXS can
  MAIN-V speak
  TENSE PRES
  OBJ (NP HEAD Spanish
        SEM (LANG ID s1)
        REF ls1)

  SEM (CAPABLE TENSE PRES
        AGENT h1
        THEME (SPEAK AGENT h1
                 THEME s1))

  REF (ABLE-STATE AGENT Suzanne
        ACTION (USE-LANGUAGE AGENT Suzanne
                  LANG ls1)))
```

The outermost category is the syntactic category, sentence. It has many ordinary syntactic features, subject, object, and verbs. The subject is a noun phrase that describes a human and refers to a person named Suzanne, the object a language, Spanish. The semantic structure concerns the capability of the person to speak a language. In the knowledge base, this becomes Suzanne's ability to use Spanish as a language.

## 2.2. Evidence for Interpretations

Our task is to model how a hearer could possibly recognize the speech act attempted by the speaker. The utterance provides certain clues to the hearer, but we have already seen that utilizing them may be complex. Our approach is a type

of pattern matching in which patterns of linguistic features that match the utterance each select a range of possible partial speech act interpretations. The output of the various rules is combined by unification at each level of the parse tree, to yield a more restricted set of more complete interpretations. This method has the advantage of being very similar to other linguistic computations, in that it is incremental, can express apparently arbitrary connections between signals and their interpretations, and can be computed with the same basic engine. Another advantage is the allowance for ambiguity, which leads to a smooth interface with plan reasoning processes. In this section we will examine various patterns, introducing any extra notation for rules as we go.

Rules consist of a set of features on the left-hand side, and a disjunction of partial speech act descriptions on the other. A rule should be interpreted as saying that any structure matching the left hand side *must* be interpreted as one of the speech acts indicated on the right hand side. The speech act descriptions themselves are also in slot/filler notation, as they are knowledge base entities. Their categories are simply their types in the knowledge base's action abstraction hierarchy, in which the category SPEECH-ACT abstracts all speech act types. Slot names and filler types also are defined by the abstraction hierarchy, but a given rule need not specify all slot values. Many of the phenomena that we mention here will be examined thoroughly at the end of the chapter.

Here is a lexical rule: the adverb "please" occurring in any syntactic unit signals a request, command, or other act in the directive class.







A large class of sentential adverbs is associated primarily with Inform acts.

- (31) a: Clearly she's our best candidate.  
       b: The cover was intact, fortunately.  
       c: They're evidently quite hot.

They are used to convey the speaker's attitude or degree of belief in the content of the sentence, and are able to enforce an Inform interpretation.

- (32) a: Actually, I'm pleased to see you.  
       b: Surprisingly, I'm leaving next week.  
       c: \*Unfortunately, I promise to obey orders.

Sentence (a) isn't quite a greeting, although it would most likely be one without the adverb. In (c) the adverb clashes with the explicit performative Promise. To some extent this is due to the inconsistency of the adverb's attitude and a sincere Promise, but usually it is not possible to comment on an explicit performative as you do it. Exceptions occur for adverbs whose semantics are highly appropriate to the act ("We proudly announce...."), and for structures in which the act itself is an infinitive complement.

A number of useful generalizations are based on sentence type. All previous work has emphasized that declarative sentences are assertions (when they are not explicit performative utterances!), imperative sentences are requests or commands, and yes/no questions are questions (or Requests for Inform acts, an analysis we will discuss eventually.) Ignoring the vestigial indicative/subjunctive distinction in English, we could refer to sentence type as MOOD, with possible values DECL(arative), IMPER(ative), YES-NO-Q, and WH-Q. Rules to handle these

cases need to allow for the possibility of other interpretations, since although these interpretations are common, exceptions are too.

(S MOOD DECL)      =(5)=>

((INFORM-ACT PROP V(REF))  
(SPEECH-ACT))

(S MOOD IMPER)      =(6)=>

((COMMAND-ACT ACTION V(REF))  
(SPEECH-ACT))

(S MOOD YES-NO-Q)      =(7)=>

((ASK-Y/N-ACT PROP V(REF))  
(SPEECH-ACT))

(S MOOD WH-Q)      =(8)=>

((ASK-WH-ACT DESCRIPTION V(REF)  
QUERY-TERM V(REF WH-QUERY))  
(SPEECH-ACT))

The value function V returns the value of the specified slot of the sentence. Thus our declarative rule has the proposition slot PROP filled with the value of the REF slot of the whole sentence. The literal meaning of the sentence is exactly the proposition that the speaker is informing the hearer of. Our innovation here is that, since the rule serves to constrain the range of possible interpretations, it must allow for the other uses of declarative utterances (explicit performatives, for instance.) Therefore the right-hand side suggests the Inform interpretation but also includes a very abstract (or generic) SPEECH-ACT. In the notation we have instead of a single category/slot/filler structure a list of such structures as possible interpretations.

The imperative rule, analogously, specifies that the action being commanded is exactly the content of the utterance, and that there may be alternative interpretations. The rule for yes/no questions is very similar. The WH rule assumes that the syntactic structure dominated by the WH word can be found in a top-level slot called WH-QUERY, as in [Allen 87]. The speech act corresponding to WH question contains slots both for the entire proposition describing the variable embedded in it, and for that variable explicitly. This allows the description to make use of other, non-queried variables. We will treat sentence types and the MOOD feature in extensive detail later.

Mood often figures in more specific patterns. Interrogative sentences with modal verbs and a subject "you" are typically requests, but may be some other act:

```
(S MOOD YES-NO-Q
  VOICE ACT
  SUBJ (NP PRO you)
  AUXS {can could will would might}
  MAIN-V +action)          =(9)=>

  ((REQUEST-ACT ACTION V(REF ACTION))
   (SPEECH-ACT))
```

This rule interprets "Can you...?" questions as requests, looking for the subject "you" and any of these modal verbs. In this rule, the value function V follows a chain of slots to find a value. Thus V(REF ACTION) takes the value REF slot and pulls out the value of the ACTION slot.<sup>1</sup>

---

<sup>1</sup>This order is reversed from [Allen 87] to correspond with the intuitions of most readers.

Some rules are based in the semantic level. For example, the presence of a benefactive case may mark a request or offer, or it may simply occur in a statement or question.

```
(S MAIN-V +action
  SEM (? BENEF ?))      =(10)=>
  ((DIRECTIVE-ACT ACT V(REF))
   (OFFER ACT V(REF))
   (SPEECH-ACT))
```

Recall that we distinguish the semantic level from the reference level, inasmuch as the semantic level is simplified by a strong theory of thematic roles, or cases, a small standard set of which may prove adequate to explain verb subcategorization phenomena [Jackendoff 72] The reference level, by contrast, is the language of the knowledge base, in which very specific domain roles are possible. To the extent that referents can be identified in the knowledge base (often as skolem functions) they appear at the reference level. This rule says that any way of stating a desire may be a request for the desideratum of the want<sup>2</sup>.

```
(S MOOD DECL =(11)=>
  VOICE ACT
  TENSE PRES
  REF (WANT-ACT ACTOR !s))
  (REQUEST-ACT ACT V(REF WANT-ACT DESID))
```

---

<sup>2</sup>A case can be made for Wanting as a voluntary action or state, when it is used as here in the sense of intention. When it encompasses desires which the agent has no intention of acting on, it no longer has any element of will or action. The interested reader is referred to [Cohen 86] for more sophisticated intention operators. The distinction between actions and states will remain a problem for knowledge representation for some time to come.

(SPEECH-ACT))

It will match any sentence that can be interpreted as asserting a want or desire of the agent, such as

- (33) a: I need a napkin.  
b: I would like three pounds of barley and some garlic.

The object of the request is the WANT-ACT's desideratum. (The desideratum is already filled by reference processing.) One may prefer an account that handles generalizations from the REF level by plan reasoning; we will discuss this point later. For now, it is sufficient to note that rules of this type are capable of representing the conventions of language use that we are after.

### 2.3. Applying the Rules

We now consider in detail how to apply the rules. A summary of their properties appears below.

**RULE: LHS => RHS**  
**LHS: ( CAT <SLOT FILLER>\* )**  
**CAT: ID**  
**SLOT: ID**  
**FILLER: ID | LHS | WORD | LIST | VALUE-FN**  
**RHS: ( LHS\* )**  
**LIST: ( FILLER+ )**  
**WORD: ID**  
**ID: a string of one or more alphanumeric characters, including -**  
**VALUE-FN: V (ID+ )**

**VALUE-FN** is a function returning the value of a specified slot from the left-hand side of a rule, used only on the right side.  
 a **WORD** must be an English word  
 syntactic categories may have feature slots, word slots, category slots.  
**SEM** and **REF** are category slots.  
**SEM** categories' slots are the 10 or so thematic roles

REF categories' slots are the corresponding knowledge base roles

RHS's are all RIF speech acts.

sometimes we will replace a LHS structure with its name, for readability, as Suzanne for some complex database entity.

#### Unification of LHS's

--categories must match

? matches any category

SEM and REF categories have abstraction hierarchies,

so a type unifies restrictively with any type abstracting it.

--if a slot is present in both, the values must unify.

a word unifies with a list if it is a member of the list, to yield the word. IDs must match exactly.

--a slot present only in one LHS appears in the final result.

For now, assume that the utterance is completely parsed and semantically interpreted, unambiguously, like the sentence "Can you speak Spanish?" as it appeared in Sect. 2.1. We repeat it here for convenience

```
(S MOOD YES-NO-Q
  VOICE ACT
  SUBJ (NP HEAD you
        SEM (HUMAN ID h1)
        REF Suzanne)
  AUXS can
  MAIN-V speak
  TENSE PRES
  OBJ (NP HEAD Spanish
        SEM (LANG ID s1)
        REF ls1)

  SEM (CAPABLE TENSE PRES
       AGENT h1
       THEME (SPEAK AGENT h1
                THEME s1))

  REF (ABLE-STATE AGENT Suzanne
       ACTION (USE-LANGUAGE AGENT Suzanne
                LANG ls1)))
```



Interpretation of this sentence begins by finding rules that match with it. The matching algorithm is a pattern matcher in the same spirit as a standard unification or graph matcher. It requires that the category in the rule match the category in the input. All slots present in the rule must be found on the category, and have equal values, and so on recursively. Slots not present in the rule are ignored. If the rule matches, the structures on the right hand side are filled out and become partial interpretations.

For example, consider the simple rule given earlier for yes/no questions acting as requests.

```
(S MOOD YES-NO-Q
  VOICE ACT
  SUBJ (NP HEAD you)
  AUXS {can could will would might}
  MAIN-V +action) ==>
  ((REQUEST-ACT ACTION V(ACTION REF))
   (SPEECH-ACT))
```

It requires the outermost syntactic category S, for sentence, which the Spanish sentence has. The first slot, MOOD, has the value YES-NO-Q, and indeed the sentence has a MOOD slot with this value. Likewise the next slot in the rule, VOICE, has a corresponding slot in the sentence with its value, ACT(ive.) The next slot, SUBJ, has an embedded structure which must be descended recursively. The embedded structure has the category NP, as does the filler of the sentence's SUBJ slot, and the rule's one HEAD slot does appear with filler "you" in the sentence. The other slots on the sentence's NP are ignored. Now, the rule asks for

the auxiliary verb slot AUXS to have one of a list of values; the sentence AUXS happens to have the first of these, "can". The rule requires the main verb to be of type +action; "speak" in the sentence is so marked in its lexical entry.

There are two partial interpretations generated by this rule. The second, SPEECH-ACT, requires no elaboration at this point. The first, the REQUEST-ACT, needs to fill in the action being requested. The value function V specifies the slots in the sentence to descend, taking the contents of the sentence REF's ACTION slot. The requested action is thus the action that is described by the embedded clause. Here is the set of two interpretations:

```
((REQUEST-ACT ACTION (USE-LANGUAGE AGENT Suzanne
                        LANG ls1)))
(SPEECH-ACT))
```

The mood rule for yes/no questions, reproduced here,

```
(S MOOD YES-NO-Q)      =(12)=>
  ((ASK-ACT PROP V(REF))
   (SPEECH-ACT))
```

produces a set of two interpretations:

```
((ASK-ACT PROP(ABLE-STATE AGENT Suzanne
                ACTION (USE-LANGUAGE AGENT Suzanne
                        LANG ls1)))
(SPEECH-ACT))
```

We need a few general rules to fill in information about the conversation:

(?) =(13)=> ((SPEECH-ACT AGENT !s))

This rule says that an utterance of any syntactic category maps to a speech act with agent specified by the global variable !s. (Speaker and hearer are assumed to be contextually defined.) The partial interpretation it yields for the Spanish sentence is a speech act with agent Mrs. de Prado:

((SPEECH-ACT AGENT Mrs. de Prado))

This is the third set of interpretations. The second rule is analogous, filling in the hearer.

(?) =(14)=> ((SPEECH-ACT HEARER !h))

For our example sentence, it yields a speech act with hearer Suzanne.

((SPEECH-ACT HEARER Suzanne))

We now have four sets of partial descriptions, which must be merged.

## 2.4. Combining Partial Descriptions

The combining operation can be thought of as taking the cross product of the sets, merging partial interpretations within each resulting set, and returning those combinations that are consistent internally. Thus, since we interpret the right hand

side of each rule as a disjunction and the set of matching rules as a conjunction, the resulting list of interpretations is a disjunction (OR not XOR), of which multiple interpretations may apply.

The operation of merging partial interpretations is actual unification or graph matching; when the operation succeeds the result contains all the information from the contributing partial interpretations. Above, we had four sets of partial interpretations. The cross product of our first two sets is simple; it is the pair consisting of the interpretation for speaker and hearer. These two can be merged to form a set containing the single speech act with speaker Mrs. de Prado and hearer Suzanne. The cross product of this with the results of the mood rule contains two pairs. Within the first pair, the ASK-ACT is a subtype of SPEECH-ACT and therefore matches, resulting in a request with the proper speaker and hearer. The second pair results in no new information, just the SPEECH-ACT with speaker and hearer. (Recall that the mood rule must allow for other interpretations of yes/no questions, and here we simply propagate that fact.)

Now we must take the cross product of two sets of two interpretations, yielding four pairs. One pair is inconsistent because REQUEST-ACT and ASK-ACT do not unify. The REQUEST-ACT gets speaker and hearer by merging with the SPEECH-ACT, and the ASK-ACT slides through by merging with the other SPEECH-ACT. Likewise the two SPEECH-ACTs match, so in the end we have an ASK-ACT, REQUEST-ACT, and the simple SPEECH-ACT.

((REQUEST-ACT AGENT Mrs. de Prado  
 HEARER Suzanne  
 ACTION (USE-LANGUAGE AGENT Suzanne  
 LANG ls1)))

((ASK-ACT AGENT Mrs. de Prado  
 HEARER Suzanne  
 PROP (ABLE-STATE AGENT Suzanne  
 ACTION (USE AGENT Suzanne  
 OBJECT ls1)))

((SPEECH-ACT AGENT Mrs. de Prado)  
 HEARER Suzanne))

At this stage, the utterance is ambiguous among these interpretations. Consider their classifications in the speech act hierarchy. The third abstracts the other two, and signals that there may be other possibilities, which it also abstracts. Its significance is that it allows the plan reasoner to suggest such further interpretations, and it will be discussed later. If there are any expectations generated by top-down plan recognition mechanisms, say, the answer in a question/answer pair, they can be merged in here.

## 2.5. Discussion

We have used a set of incremental rules to build up multiple interpretations of an utterance, based on linguistic features. They can incorporate lexical, syntactic, semantic and referential distinctions. To demonstrate the effectiveness of the system, we will consider what happens to "Can you speak Spanish, please?"

((S MOOD YES-NO-Q  
 VOICE ACT  
 SUBJ (NP HEAD you  
 SEM (HUMAN ID h1)  
 REF Suzanne)

AUXS can  
 MAIN-V speak  
 TENSE PRES  
 OBJ (NP HEAD Spanish  
       SEM (LANG ID s1)  
       REF ls1)  
 ADV please  
  
 SEM (CAPABLE TENSE PRES  
       AGENT h1  
       THEME (SPEAK AGENT h1  
               THEME s1))  
  
 REF (ABLE-STATE AGENT Suzanne  
       ACTION (USE-LANGUAGE AGENT Suzanne  
               LANG ls1)))

The only difference between this sentence and the previous example is that this one includes the adverb *please*. The word has no corresponding linguistic presence in the logical form, nor in the knowledge representation. The rules that match it are the same as before, with the addition of the "please" rule.

(? ADV please) ==>  
 ((DIRECTIVE-ACT))

This rule matches its wildcard category against the S of the sentence, and finds the adverb slot with *please* in it. The resulting interpretation is simply

((DIRECTIVE-ACT))

Thus, the complete set of partial interpretations for the sentence is

((DIRECTIVE-ACT))  
 ((SPEECH-ACT AGENT Mrs. de Prado))  
 ((SPEECH-ACT HEARER Suzanne))

((REQUEST-ACT ACTION (USE-LANGUAGE AGENT Suzanne  
LANG ls1)))  
(SPEECH-ACT))

((ASK-ACT PROP(ABLE-STATE AGENT Suzanne  
ACTION (USE-LANGUAGE AGENT Suzanne  
LANG ls1)))  
(SPEECH-ACT))

The cross product of the first three sets, with merging, is

(DIRECTIVE-ACT AGENT Mrs. de Prado  
HEARER Suzanne)

since a directive act is a specialization of a speech act. The cross product with the next set yields two interpretations, the request specializing the directive act and the directive act specializing the generic speech act.

((REQUEST-ACT AGENT Mrs. de Prado  
HEARER Suzanne  
ACTION (USE-LANGUAGE AGENT Suzanne  
LANG ls1)))

DIRECTIVE-ACT AGENT Mrs. de Prado  
HEARER Suzanne)

The final cross product has the same result, because the SPEECH-ACT merges with both interpretations, but the ASK-ACT merges with neither. Thus, the "please" rule constrains the results of merging to be directive acts.

The Spanish example demonstrates the power of the word *please*, which overrides our preference for a yes/no interpretation. But we should also explain why the yes/no interpretation is preferred in the unmarked case. One possible explanation

is that, for sentences taken out of context, sheer frequency of use plays a role in our intuitions. Americans are simply never asked to speak a particular language. But there is also a linguistic-semantic reason for this. The sense of "speak" used in English for language fluency has no role for the utterance content. "Can you speak Spanish?" and "Can you read Spanish?" are not specific enough to indicate what is to be said or read, and are therefore inadequate to express most requests for use of a language. This additional information would need to be very obvious in context, or be specified in an additional utterance:

(34) Can you read Spanish? This paper's important, but it's in Spanish.

This request is spread over two sentences, and we therefore cannot justify labelling the first sentence a Request act on its own. The lone sentence is almost impossible to recognize as a request.

It is clear that some cues are much stronger than others. We have incorporated this distinction in a very simple way: a sufficiently strong cue has only one possible interpretation, while weaker cues leave the range of alternatives open. Even so, we sense that one interpretation from the right hand side may be favored in some rules, and that some possible interpretations are extremely unlikely. For fine-tuning of the model, we might be able to add to each of the right-hand interpretations a weight, which is derived from frequency data and would for a human incorporate social class, idiolect, and so on. Each pattern is then evidence for a distribution of interpretations. This does not affect our central claim, which



is that the evidence combines incrementally to constrain the range of interpretations. But it is probably necessary for any system with broad coverage.

### 2.5.1. Another Example

Explicit performative utterances [Austin 62] deserve special mention. First, they have very distinctive surface form. Only they may contain the word "hereby". They are also declarative, active, utterances whose main verb identifies the action explicitly. Second, they have very simple interpretations. The sentence meaning corresponds exactly to the action performed. (There is a remote chance that the sentence has a habitual reading, but we will ignore it here.)

```
(S MOOD DECL
  VOICE ACT      =(15)=>   (V(REF))
  MAIN-V +performative
  TENSE PRES)
```

One might be tempted to insist that the subject must be "I", but there are other acceptable forms:

- (35) a: We proudly introduce Admiral Grace Hopper.  
       b: The Society for Women Engineers proudly introduces Admiral Grace Hopper.

Let us see how such a sentence is processed. It might be represented like this:

```
(S MOOD DECL
  VOICE ACT
  SUBJ (NP HEAD (PRO WORD I
                  NUM 1s)
        SEM (HUMAN ID hu1
              NUM sing)
        REF Jane)
```

ADV proudly  
 MAIN-V introduce  
 TENSE PRES  
 OBJ (NP PREMODS (Admiral Grace)  
     HEAD Hopper  
     SEM (HUMAN ID h1  
         NAME Grace Hopper  
         TITLE Admiral)  
     REF AGH)

SEM (INTRODUCE TENSE PRES  
     AGENT hu1  
     THEME h1)

REF (INTRODUCE-ACT AGENT Jane  
     PARTY1 AGH))

This is a declarative, active sentence with a performative main verb, and the subject is first person singular. The performative rule, number 15, matches. We have the partial interpretation:

(INTRODUCE-ACT AGENT Jane  
     PARTY1 AGH))

Another rule that matches is the declarative rule:

(S MOOD DECL)      ==>  
     ((INFORM-ACT PROP V(REF))  
     (SPEECH-ACT))

It yields two partial interpretations:

((INFORM-ACT PROP INTRODUCE-ACT AGENT Jane  
                                     PARTY1 AGH))  
 (SPEECH-ACT))

From the context rules,

(?) ==> ((SPEECH-ACT AGENT !s))  
 (?) ==> ((SPEECH-ACT HEARER !h))

we get expressions for the speaker and audience.

((SPEECH-ACT AGENT Jane))  
 ((SPEECH-ACT HEARER aud9878))

We have again generated four sets of partial interpretations. The lone Introduce act, from the explicit performative rule, unifies only with the generic speech act of the declarative rule, and thus eliminates the Inform act. The other two sets add the speaker and hearer. This is the result:

(INTRODUCE-ACT AGENT Jane  
                   PARTY1 AGH  
                   HEARER1 aud9878))

The other role PARTY2 should become filled from the hearer role. Our sense that the utterance is a statement comes from the fact that of course declarative utterances *are* prototypically Informs, and this is reflected in our interpretation process.

The word "hereby" cues a performative in the same way as "please" cues requests, and even more strongly so:

(? ADV hereby) =(16)=> V(REF)

There are passive voice performatives, not captured by the performative rule, that

are odd without "hereby". "You are hereby informed that homeowners must have chimney filters", which would be treated as an Inform of an Inform by the declarative rule, is constrained to a simple Inform act by the "hereby" rule.

Having demonstrated the basic mechanism that generates speech act interpretations, we will in chapter 3 look at some linguistic cues in further detail. There we will also examine some related linguistic issues for which our method has implications. Later chapters will establish the role of plan reasoning in speech act interpretation.

### 3. Linguistic Constraints II: Case Studies and Limits

The linguistic patterns in Chapter Two were presented simply, in order to focus on the techniques for speech act interpretation. We now examine these patterns in detail, demonstrating empirically that they serve as pragmatic signals. This leads to a refined speech act hierarchy, as well as to the limitations of linguistic cues and our understanding of them.

The linguistic patterns that we have studied in detail include sentence type, and the lexical items *hereby* and adverbial *please*.

#### 3.1. Sentence Type

Sentence type or mood has always been assumed to play a prominent role in speech act interpretation. This role has been overstated and oversimplified at times, and even with very broad observations we can refine this traditional view significantly.

In the absence of other indicators, sentence type provides a rough guide to speech act type. Earlier we made reference to sentence mood, but this four-way distinction needs refinement to gain coverage of the majority of English utterances. In addition to complications of the main sentence types in English, there is a wide variety of minor types, which are ordinary enough but simply less frequent than the ones mentioned so far. Isolated noun phrases may serve as question answers, with falling intonation. They may be questions, requests, or offers, with rising intonation, or exclamations, with contrastive stress. Thus noun phrases are

acceptable forms under many circumstances.

- (36) a: Jam.  
       b: Jam?  
       c: Jam?!

There are several types of alternative questions, which spell out the possible answers for the hearer:

- (37) a: Would you like coffee, tea, or cocoa?  
       b: Are you coming or not?  
       c: What would you like to drink? Coffee, tea, or cocoa?

They may resemble yes/no questions but provide a disjunction of values that would be appropriate for a wh-question (a). They may resemble yes/no questions and specify the disjunction of a positive and negative value (b). They may also take the form of a list alone, possibly preceded by a wh-question (c).

Sentences with question form and emphatic falling intonation may act as an exclamatory assertion; this may occur with a negative form, or with stress on the verb and subject:

- (38) a: Hasn't she grown!  
       b: Has she grown!

Sentences with statement form and rising intonation may act as yes/no or wh-questions.

- (39) a: You're leaving town on Thursday?  
       b: You're leaving town when?

We will confine our discussion to these, although there is a simply wonderful assortment of sentence forms which are somewhat less common [Leech '75]. Many of these

are special cases of those above, such as biased questions, questions with more than one wh-word, tag questions, echo questions, wh-echoes, and reported and short forms of all of these. Others, such as vocatives and forms specific to greetings & other social actions, are not. A final category, consisting of backchanneling and attention signals, is arguably not a category of sentence forms at all though such forms clearly have a role in communication.

The sentence types we have identified provide a mapping from surface features to speech act types for which they are suggestive evidence. The MOOD feature that we used earlier is a composite of several syntactic features. A declarative sentence, for instance, has a subject followed by a verb phrase with any of several forms. A yes-no question has a subject and verb phrase, but shows subject/auxiliary inversion. A wh-question is the same but with a fronted wh-term (who, how, etc.) Imperatives have an imperative verb form and the subject is often implicit. The table below summarizes these features of sentence type.

type	subj/aux inv	subject	special
declarative	-	+	-
imperative	-	-/you/someone	imperative verb form
y/n question	+	+	-
wh question	+	+	fronted wh-term

This set of features would be an adequate basis for our earlier sentence form rules, and will serve as a definition for the shorthand of MOOD values. The extended version, with our new sentence types, is shown with a very simple intonation summary. In English, a final rise in intonation suggests incompleteness

(questioning), while a final fall conveys certainty. Thus it is possible to distinguish spoken questions in declarative form from statements, for example. One would like to think that punctuation would play the same role in informal written English, but people appear to be inconsistent (see for example the data in [Allen 89].) The syntactic data here are uncontroversial; use of intonation for distinguishing actual speech acts is a much more difficult question.

form	subj/aux inv	subject	intonation	special	SA type
declarative	-	+	f		Inform
declar. y/n	-	+	r		Ask-YN
declar. wh	-	+	r	wh-term	Ask-WH
y/n question	+	+	r		Ask-YN
y/n rhetorical	+	+	r	rhet tone	Ask-Rhet
y/n exclamat.	+	+	f		
y/n alternate	+	+	f	disjunct NP's	Ask-WH
wh question	+/-	+	f	fronted wh-term	Ask-WH
wh rhetorical	+/-	+	r	fronted wh-term	Ask-Rhet
imperative	-	-/you/someone	f	imperative verb	Directive
NP statement	N/A	-	f		Inform
NP q/r	N/A	-	r		Inform
interjections	N/A	-	var.	fixed forms	various

Many different speech act types occur with each of these values, but in the absence of other evidence an utterance with the given features is likely to have the corresponding speech act type. Provisional definitions for the speech act types are given in Section 5.2, with the exception of rhetorical questions.

Note that unlike previous work, we do not treat questions as Requests to Inform. The logical conditions on these acts are not very different, but there are several language-based reasons for making the distinction. First, English embodies the



distinction in its fundamental sentence types, even though these types are used for various purposes. People are able to reason about questions *per se*. Second, most languages of the world distinguish questions and requests in their fundamental sentence types [Sadock *ng*], suggesting that this is a very useful cognitive distinction. Third, "please" is not commonly regarded as acceptable with questions, although there are requests to inform [Sadock 74]. Therefore we will use distinct speech act types for question classes, as well as requests.

### 3.2. Hereby

Certain sentential adverbs are firmly associated with certain speech act types. The adverb *hereby* is regarded in the speech act literature as a marker of, and even a test for, explicit performative utterances. As one would expect, it is derived from the adverb of place, *here*, and the preposition *by*. Archaically it meant "near this place" or less spatially "in this connection". The Oxford English Dictionary's only extant meaning is "By, through, or from this fact or circumstance; as a result of this; by this means." It can still be used as a referring expression, before the main verb or in final position:

(40) She called him a cad. He was humiliated hereby, but said nothing.

To substantiate the role of *hereby* as evidence for an explicit performative utterance, we searched some 42 million words of text from the Associated Press Wire Service. Ken Church of AT&T Bell Labs kindly provided the expertise and the stemming algorithm. There were 52 occurrences of the word, from which we

removed 18 duplicate quotations by hand. Of the remaining 34, 27 are clear explicit performatives declaring, proclaiming, announcing, and so on. The final seven are as follows. One appears in a document, a bid to buy a hotel chain.

- (41) ...we have available sufficient funds to consummate the transactions contemplated hereby.

This is the only occurrence in final position; the rest precede the main verb. The transactions were contemplated in earlier sentences, but of the same document. So one could understand it as referring to the entire document as an explicit performative action. This ballooning of speech acts into larger chunks of text is a phenomenon unaccounted for by current theories. The remaining six utterances were produced by non-native speakers of English.

- (42) I am hereby announcing a proposal which I am addressing to....

This one is unique in being in a progressive tense, and has an analysis similar to the last example. It comes from Poland. The others are from the Middle East.

- (43) a: Hijacker: We hereby re-announce our refuelling request....  
 b: We hereby make it clear that we do not have the slightest intention....  
 c: Khomeini: I hereby want all the dear people ... to be patient....  
 d: Kidnappers: ...we hereby enclose with this statement the recorded message ....  
 e: I am hereby the deputy foreign minister of Iran  
 officially declaring that there is no obstacle....

All of these cases deviate at least slightly from our use of the word. In this dialect of international rhetoric, it appears to mean something like "officially", just as *please* becomes a way for some non-native speakers to express honorifics. The verbs with which it appears are unusual ones. "Re-announce" is simply novel;

"make it clear" isn't quite performative inasmuch as the agent ultimately can't ensure this effect. "Want" isn't performative if taken literally, but it can be construed to have the sense of a request. It would be particularly helpful to know the equivalent expression for this case, in the speaker's first language. In (d) "enclose" suggests the spatial sense of *hereby*, but the full quote suggests the "officially" reading as well.

- (44) On the occasion of Terry Anderson's birthday and in response to your letters, and according to his desire to send you a recorded message, we hereby enclose with this statement the recorded message on video tape. (the kidnappers said.)

In (e) one might be tempted to understand the adverb as displaced from the verb, along the lines of the explicit performative

- (45) ...PLO, hereby once more declare that I condemn terrorism....

But it occurs within the noun phrase, and there is already a preverbal adverb, so the speaker appears to be emphasizing (surely not self-appointing!) his office.

If we wish to handle the full range of these quotations, including utterances of non-native speakers, we will need to treat *hereby* as lexically ambiguous among the pure performative sense, the generalized "officially" sense, and possibly the spatial sense. The pure performative sense was seen here 60% of the time; 100% if non-native speakers are excluded.

### 3.3. Please

In American and British English, *please* is most often used with polite requests and commands, although it can also be a transitive verb or appear in other idiomatic expressions. The primary role of the polite word is pragmatic, we claim, rather than syntactic or classically semantic. We will examine its uses in detail, then show what information a discourse system must have about it, and finally how to use it in computing speech act interpretations.

Most dictionaries classify *please* as a verb, intransitive and transitive, and note that it can be used 'for politeness'. In actual usage *please* is most commonly an adverb, as we shall demonstrate.

#### 3.3.1. A Little Etymology

For our purposes, there are three senses of *please*. First in most dictionaries is the common transitive verb meaning to gratify. It originally took a dative *with, to*, etc, but the case is no longer marked and is regarded as accusative (a). It can occur with a formal subject only, and a complement (b), with a reflexive (c), passive (d), or with other more minor variations.

- (46) a: Congress never quite pleases voters.  
       b: It pleases her to destroy things.  
       c: Cats please themselves.  
       d: We're so pleased to see you.  
       e: He puts on his palette the things that please most palates. (AP)

The impersonal form above in (b) was once used in a number of deferential expressions, in this same sense of volition or desire. They behave as adverbial

phrases:

- |                                     |                           |
|-------------------------------------|---------------------------|
| (47) a: {and, an, if} it please you | [, so eat we now.]        |
| b: {may it, will it} please you     | [now to eat?]             |
| c: so please you                    | [, the guests have come.] |
| d: please your honor                | [, it's the truth.]       |
| e: {may it} please God              | [ he comes home safe.]    |
| f: please you                       | [lend me your horse.]     |

The form (f) is the shortest form that occurs in Shakespeare, so that the form we are interested in is more recent. One last example of the transitive verb is this reflexive imperative:

- (48) Please yourself, then!

It is possible to use the imperative very politely, but here it's sarcastic: since you won't listen to me, I'll just direct you to do what you will anyway. An intransitive counterpart is closely related to the transitive form. Compare

- (49) a: We aim to please.  
b: We aim to please customers.

There is another sense of the verb which has the meaning reversed. Most people find it unacceptable in general (a below), but it occurs in many common phrases (b-d) in a *wh*-extracted form.

- (50) a: Cats please to lie in the sun.  
b: ... right to associate with whom they please. (AP)  
c: ... a right in 1988 to worship where we please. (AP)  
d: ... brain surgeons can live however they please and still ... (AP)

It has a transitive form which is obsolete. The reversed form of *please* appeared suddenly in the early 15th century; the OED has this to say about it:

The history of this inverted use of *please* (observed first in Scottish writers) is obscure. But exactly the same change took place in the 14th c. in the use of the synonymous verb *Like*, where the impersonal "it liked him", "him liked", became "he liked" ca. 1430. It may therefore be assumed that "I please" was similarly substituted for "it pleases me", "me pleases" (c. 1440.) ....The remarkable thing in the case of *please* is that the sense was already logically expressed by the passive....

This seems to be a consequence of the same development in English that gave us the modal verbs (see [Lightfoot 79], for example.), namely the switch from SOV to SVO word order, which occurred abruptly ca. 1500. In any case, the OED suggests that the optative *Please!* originally derived from the adverbial phrase *please you*. However, it adds that we now analyze it as an imperative of the flipped verb or a reduction of the flipped "if you please". This form is reinforced by contact with French, where the *you* is (an unmarked) dative.

We conclude from all this that the adverbial sense is today well removed from the primary transitive verb, and its semantics cannot be taken directly from there. It also seems unlikely that six-year-olds who are taught to "say please" appreciate the connection with "as you please". For semantic purposes, the adverb is best allowed to stand on its own.

### 3.3.2. Uses of *Please*

Now let's consider how the adverb is used. We will confirm the adverbial view in examining a large body of data, roughly a year's worth of Associated Press wire service text (AP). But first, we summarize what we already know about it.

Sadock [Sadock 74] discusses the adverbial use at some length. As an adverb,

*please* is a sentential one with three common occurrences: sentence-initially, sentence-finally separated by a pause, or internally, preceeding the main verb.

- (51) a: Please tell your fellow soldiers we are thinking of them. (AP)  
 b: Help us, please! (AP)  
 c: May I please have your autograph? (AP)

A number of variations on this theme can be observed. In (a below) the entire clause appears in apposition. In (b) and (c), initial *please* is separated from the main verb by a vocative expression. In (d) and (e) it occurs among the modifiers of the verb phrase, where a vocative expression could also occur.

- (52) a: - and, please note, that takes time. (AP)  
 b: Please, Sir, can I have some more? (Dickens, *Oliver Twist*)  
 c: it's like, please, someone shoot me if I ever say that. (AP)  
 d: Come in the office, please, with your children. (AP)  
 e: Would you identify for us then, please, three specific programs ...(AP)

In each of these cases, adverbial *please* has been associated with a polite directive act, and the act desired is given by the main verb of the sentence. The (a) request is about the discourse, and the (c) request is jocular or rhetorical in tone, but both are requests. Sentence (d) may be a polite command, also a directive. Here are two other requests:

- (53) a: Take my deli -- please! (AP quoting Henny Youngman)  
 b: Diane, the diamonds please. (AP)

The humor in (a) is a pragmatic pun. It uses *please* to take an idiomatic topicalization and re-interpret it as a literal, outrageous request. The (b) utterance is typical of another form of directive, a noun phrase. A pause is needed between

the noun phrase and adverb, as with full sentences. With the possible exception of a few idiomatic cases, *please* serves as an indicator of a polite directive. In this role the word has no classical semantics, but very clear pragmatics which we must utilize.

Other cases are very specific. Accepting an offer politely is one. Paraphrases with the same use are also given. Our intuition is that these expressions are short for repeating the entire offer as a request.

- (54) a: Yes, please [do wrap my package.]  
       b: Yes, I would [like some tea.]  
       c: Yes, thank you.  
       d: Yes, please do [drop by sometime.]

Another is a request for attention, as in a restaurant.

- (55) a: Please, Miss...  
       b: Excuse me....  
       c: Waiter!

A third, with heavy stress, rudely discredits the previous speaker. Its sarcasm does not invert its directive sense but its politeness. It would be interesting to see if this leads to a view of sarcasm consistent with extensive data.

- (56) a: Oh, please!  
       b: Spare us!  
       c: Oh, come off it!  
       d: Oh, cut the nonsense!  
       e: Oh, gimme a break!

lp All uses of *please* we have seen so far have been directive acts. With a full probability theory, we could write a rule that expresses the likelihoods of the



possible specializations.

### 3.3.3. The Data

One might gather from a dictionary definition that the transitive verb is the most common, followed by the intransitive. In 42 million words of Associated Press wire service text we found that, although the verb sense was about four times as common as the adverb, the uninflected form was four times as likely to be an adverb as the verb. The gross breakdown is shown in the table below.

620	please
1	Please-Some
4	please-raise-my-taxes
1	hard-to-please
31	pleases
47	pleasing
1	audience-pleasing
6	crowd-pleasing
1226	pleased
-----	
1863	

There were 1863 occurrences of *please* and its verb forms. Of those, 620 are uninflected *please*. One difficulty of the AP data is that some occurrences of a given form are actually multiple citations of one original quote. We have attempted to eliminate duplicates for the subsequent analysis. News reportage is a hardly a domain of choice for discourse study, since it generally does not consist of dialogue, so the dominance of the verb over the adverb is not surprising. The train station data [Horrigan 77], by contrast, contain not a single verbal *please*. A non-

empirical study of the University of Birmingham corpus tended to be more than half adverbial *please*, and included occurrences of almost all archaic or idiomatic usage one could imagine. In sum, the AP data have enough adverbial *please* to be worth investigating, if not every possible variation.

The breakdown of *please* occurrences in the AP data appears in the following table.

Form	Count	Comments
Transitive verb	107	56 w/to, 40 modal, 1 it
Intransitive	6	(We aim to please.)
please God	1	
Flipped	25	All with wh-extraction.
preverbal(imper)	297	incl at least 15 voc, 9 advp, 9 with pause
final (imper)	10	
preverbal (interr)	15	modal incl 2 reported
final (interr)	9	incl 2 "may I take your order, please?"
preNP	5	incl. 2 with vocatives
postNP	16	
isolated	6	incl 3 with vocatives
indirect requests	7	
henny youngman	6	
quoted	3	
song titles	21	"Please, Please, Please", "Please, Mr. Postman" "Will you please be quiet, please?"
duplicates&typos	79	
other	11	

The findings for verbs are not surprising. They occur mostly in speculations about whether something would or is likely to please voters, customers, or other countries. The flipped sense occurs only with wh-extraction (if it is possible to wh-extract *as*.) The imperative sentences show intermixing of *please* with vocatives

and adverbial phrases before the verb, occasionally, as well as a few final *pleases*. (Of course the preverbal and initial positions are the same for most imperative sentences, so there is no need to subdivide.) The interrogative sentences all include a modal verb with *please* before the main verb or in final position, but never initially. It may occur before or after noun phrases or stand alone, with or without vocatives. There were seven instances one might term indirect requests:

- (57) a: Please, you have to get this by such and such a time.  
 b: Please, I really want to forget about that.  
 c: ...do it in reverse order if I could, please.  
 d: Moderator: Please, please, once again you're only taking time away ....  
 e: - when their candidate speaks, so please.  
 f: You want to do me a favor please?

These cases are clearly directive, but the attachment of *please* is less easy to describe. There were six instances of the ever-popular "Take my X -- please!"

The remainder are each worth commenting on.

- (58) a: Yes, please.  
 b: Oh, please!  
 c: Time, gentlemen, please!  
 d: If I could have your attention, please...

The first two are expression we have already discussed. The third is the call to close British pubs, which is a polite if indirect directive. The fourth is reminiscent of the noun phrase class, except that it is adverbial itself. The next four are genuinely tricky cases:

- (59) a: I do ask you to please keep your hearts and minds open....  
 b: Why don't you please try this word, no comment, just this one time.  
 c: Your point has been made and we are please asking you to leave.

d: We want them to please, think about this child....

Sentence (a) would trigger the explicit performative rule, yielding a disjunction based on the ambiguity of *ask* between the question and request sense. This is of course resolved by the *please* rule. If the latter adds anything to the interpretation, it is that the verb on the right means the requested action. Sentence (b) triggers the suggestion rule with "Why don't", and this rule must allow for literal questions but need not allow for a request. Hence there may be a clash with the request interpretation. This is consistent with our sense that the sentence is odd, but there is no real pressure to resolve the thing. Sentence (c) has *please* acting like a misplaced modifier: if a *please* rule insists that the verb to the right is the requested action, this and the explicit performative interpretation will fail to unify. However, both the explicit performative interpretation and a non-specific request would be consistent with the context. In (d) our sense is that the sentence was begun as an Inform rather than a directive, but switches viewpoints midway. The system would use *please* to restrict the declarative rule's open output down to a request, with the semantic WANT rule as further evidence. Here the system does miss some subtlety.

- (60) a: For more information, please contact your local legalization office.  
 b: ...and wish to keep it confidential please leave your name ....  
 c: In the event of emergency or clarification, please contact:...

The last three above are fairly clear cases of instructions, which are indeed directives, but conditioned on the hearer's goals or on events in the world. We should devise a speech act class for helpful instructions, as well as these possibly mandatory world-conditioned ones. These are the last of the eleven utterances in

the "other" category, so all the occurrences of *please* are accounted for.

There are utterances tallied above which are pragmatically interesting. Five cases are clearly pleading:

- (61) a: But the Met said, 'Oh please, Mirella, just two performances...'  
 b: Moderator: Please, please, once again you're only taking time away ....  
 c: ...and they said, 'Please, please do it.'  
 d: ..., please, please have correct change.  
 e: ...and I thought, 'Please, God, please.'

It is impossible to draw a firm line between pleading and other adverbial uses; this will be reflected in our speech act hierarchy, where pleading is a specialization of a polite request. Some of the data we counted in the preverb and other adverbial categories probably qualifies as pleading but without the suprasegmental component we can do little to distinguish them. There is also

- (62) May I take your order, please?

It could be regarded as asking for permission rather than asking for the order, as one would assume for analogous

- (63) Can I please go out to play?

To do this we would need a more specific rule that incorporates the information that the requested action is permission for the explicit action, if the subject is the speaker. We have as yet no mechanism for giving specific information high priority.

Just a few more notes on directives. They are often negative, and often with verbs that we mightn't consider +ACTION, like understand, accept, and so on. It is worth noting that while negative requests can stand on their own, the -ACTION ones really seem to need the preverbal *please*, or some other marker.

(64) (Please) understand that I've been really busy.

### 3.3.4. Taxonomy of Speech Acts

At this point we can draw a partial taxonomy of speech acts based on what we have seen in our analysis of *please*. Such a taxonomy is incorporated directly into the system in the form of an inheritance (or IS-A) hierarchy. Classifying a given utterance in the hierarchy can produce useful information even if the utterance cannot be associated with a leaf node. We may know that an utterance is directive, for instance, without being able to distinguish whether it is a request or a command. A working taxonomy includes some very specific acts, which depend on both language and culture. In most cultures there is a need to announce one's self when arriving at a dwelling, for instance, but how this is done will depend on the kind of dwelling. You can't knock on a grass hut. In English there is a shortage of forms of address for strangers, so that getting the attention of a stranger whom we need some service from becomes a very particular act.

This taxonomy of speech acts is really one subtree of human actions, which in turn is a subtree of the agent's taxonomy of the world. The speech act subtree is dominated by the generic (most abstract) speech act (not shown.) (Links with

- commissives -- promise
  - suggest
  - offer
  - accept offer
- directives -- requests -- standard polite request -- taking up an offer
  - familiar request
  - begging&pleading
  - invocation/blessing
- directing attention -- requesting a discourse referent
  - requesting attention
- instructions --
- commands -- polite command
  - rude command
  - parental
  - military

nonlinguistic acts could be built with the aid of multiple inheritance.) The class of commissives is one of a small number of classes directly beneath the generic speech act. Commissives are acts which obligate the speaker to make something true in the world which otherwise might not be the case; promising is a paradigm example. Here we add suggestions, which advocate a course of action, and offers, which bind the speaker to an action but contingent on the hearer's wish. Accepting an offer can also be seen as advocating a course of action, but this assignment is a bit muddy.

Directive acts are attempts to get someone to do something which they might otherwise not do. These may be requests, in which the speaker relies on the good will of the hearer, or commands, in which the speaker exercises power of authority

or force over the hearer. Requests may be polite or familiar, abject, or directed at a divine being. Requesting attention is an act which need not be linguistic at all, while requesting the hearer to locate a discourse referent [Perrault 78] is a peculiarly linguistic version of directing someone's attention. Instructions could reasonably be regarded as Inform acts, since they are information that one uses contingent to one's own goals. But the information is presented in a directive way, after all, so we include them here. In an educational setting they are clearly intended to be complied with.

With commands, compliance is not optional. They may be expressed politely or be very abrupt. The distinction between requests and commands is based on this necessity, which is a context-dependent factor not dependent on linguistic cues. Thus, though in the AP data we see 385 occurrences of adverbial *please*, of which 3 (.8%) are instructions and 5 (1.2%) are pleading, the remaining bulk of directives cannot be subdivided into polite requests and commands on the basis of the text alone.

### 3.4. Syntactic Complications

There are several complications that must be addressed by a linguistic theory of speech acts. We enumerate them here as open topics. There are many speech acts that have been referred to as indirect acts, in which the explicit performative verb is embedded in a non-auxiliary verb construction. These embedded speech acts should be shown to fall out of a compositional model of speech act interpretation.



There is the question of the speech act type of two conjoined acts, as well as the constraints on such conjunction. There is the question of how to explain certain syntactic phenomena in which the speech act appears to participate, even when it is not explicit in the sentence. The only issue we address here is that of the limiting cases of conventionality.

### 3.5. The Limits of Conventionality

We do not claim that all speech acts are conventional. There are variations in convention across languages, of course, and dialects, but idiolects also vary greatly. Some people, even very cooperative ones, do not respond to many types of indirect requests. There are cases in which the generalization is obvious but only special cases seem idiomatic:

- (65) a: Got a light?  
       b: Got a dime?  
       c: Got a donut? (odd request)  
       d: Do you have the time?  
       e: Do you have a watch on?

There are other cases in which the generalization is obvious but no instance seems idiomatic. If someone is responsible for an action, asking whether it's done is as good as a request.

- (66) Did you wash the dishes?

In the next examples, there is a clear logical connection between the utterance and the requested action. We can write a rule for the surface pattern, but the rule is

useless because it cannot verify the logical connection. This must be done by plan reasoning, because it depends on world knowledge. The first sentences can request an action to which they represent preconditions, the second set, effects.

- (67) a: Is the garage open?  
       b: Did the dryer stop?  
       c: The mailman came.

- (68) a: Is the car fixed?  
       b: Is your room clean?

Plan reasoning provides an account for all of these examples, and we will use it. The fact that certain examples can be handled by either mechanism we regard as a strength of the theory: it leads to robust natural language processing systems, and explains why "Can you X?" is such a successful construction. Both mechanisms work well for such utterances, so the hearer has two ways to understand it correctly. These last examples, along with "It's cold in here", really require plan reasoning.

In our approach, there is a continuum of speech acts from very literal to very indirect. If there is a gap, it is between the most conventional acts and the ones requiring the most reasoning, and this should show clearly in psycholinguistic studies. It is certainly not between literal and nonliteral forms, and so Searle is rescued from the criticisms of Gibbs. Another datum that supports this argument is conjunction:

- (69) a: I want two hamburgers, and put mustard on them.  
       b: \*It's cold in here, and get out.

These examples are a puzzle for Gordon and Lakoff. (a) is a pair of requests (they say, an indirect request and a command), and (b) an inform and a command (they say, an indirect request and a command.) We want to say that the request conveyed by "It's cold in here" is not conventional, while "I want two hamburgers" is, and the extra effort required beyond the convention interferes with processing the conjunction.

## 4. Plan Reasoning

### 4.1. Role of Plan Reasoning

#### 4.1.1. Introduction

In the last two chapters we viewed speech acts as the output of a linguistic interpretation process. Now we will shift our perspective, viewing speech acts as the representations used by agents for planning. We can then elaborate the constraints placed on speech act recognition by what we know of general reasoning about plans.

Plan reasoning contributes in several ways to speech act recognition. First, it provides the link between speech act interpretations proposed by the linguistic mechanism, and the facts in the actual context which are relevant. This allows the system to eliminate speech act interpretations if they contradict known intentions and beliefs of the agent. Second, it elaborates and makes inferences based on the remaining interpretations. This allows the system to process non-conventional speech act interpretations. Third, it could propose interpretations of its own, when there is enough contextual information to infer what the speaker might do next. For example, plan tracking could generate the expectation that the act following a question is an Inform. Fourth, plan reasoning provides a competence theory motivating many of the conventions described in earlier chapters.

To provide a context for elaborating these points, we will survey the previous use of plans in understanding discourse. Work specifically on speech acts will illustrate the potential of the plan-based approach. Broadly speaking, the current work emphasizes the first point. We will show a strong resemblance between our inferences and some classes of conversational implicature.

#### 4.1.2. Plans and Discourse

The use of action representations for natural language semantics has a long history. The first widely-used representation for actions was the *script*, [Schank 77]. Scripts are detailed scenarios listing a series of steps in a stereotyped process. A popular example is going to a restaurant: one may make reservations, get in the car, drive to the restaurant, park, enter, be seated, order, eat, pay, and leave. If a script-based system identifies a story as a restaurant story, it can follow this series of events as it occurs in the story, even inferring steps that were not explicitly mentioned. Such a story understanding system is described in [Cullingford 86]. Scripts have also been used as a basis for question-answering systems [Lehnert 78].

Scripts are relatively inflexible and unable to incorporate descriptions of unexpected events. Subsequent work took advantage of advances in planning, allowing actions to be strung together and connections between them to be inferred. [Wilensky 83] describes story understanding from this viewpoint, listing a variety of relationships that could hold among actions and goals. [Grosz 86b]

describes a dialogue system that uses plan tracking techniques to structure the dialogue as well as to determine the referents of noun phrases. [McKeown 86] used a scriptlike representation for generating natural language texts several paragraphs in length, describing objects known to a database. [Pollack 86] investigated relaxation of the assumption that domain plans are shared by both communicators, allowing one agent to reason about the other's possible misconceptions.

[Perrault 78] was the first work which explicitly treated communication as a series of actions to be modelled, following the philosophy literature. In that vein, [Litman 85] proposed the use of "metaplans", or ways an agent could use a speech act to modify a domain plan. [Grosz 87] also made use of action representations to describe discourse structure. [Appelt 85] investigated the generation of actual text from speech act descriptions, including satisfying multiple goals in a single sentence and generating object descriptions according to an explicit planning model. All of this work would be extended by notions of how speech acts can be recognized.

#### **4.1.3. Plan Reasoning with Speech Acts**

[Perrault 80] gave an account of indirect speech acts, based on the STRIPS model of planning. Speech act types were action descriptions, which could be recognized by an inference process inverse to that of constructing plans. The process was controlled by weighted heuristic search.

The logical machinery Perrault and Allen used to model agents has several components. First, agents have all the theorems of first order predicate calculus. The belief operator  $B_A(P)$  is a modal operator that has the following properties:

$$B_A(P) \rightarrow B_A(B_A(P))$$

$$B_A(P) \wedge B_A(Q) \rightarrow B_A(P \wedge Q)$$

$$B_A(P) \vee B_A(Q) \rightarrow B_A(P \vee Q)$$

$$B_A(\neg P) \rightarrow \neg B_A(P)$$

$$(\exists x)B_A(P(x)) \rightarrow B_A((\exists x)P)$$

$$B_A(P \rightarrow Q) \wedge B_A(P) \rightarrow B_A(Q)$$

It is also closed under Modus Ponens and the axioms. The knowledge operator  $K_A(P)$  is defined as true belief, and there are two other predicates for knowing. One represents knowing whether:

$$Knowif_A(P) \Leftrightarrow K_A(P) \vee K_A(\neg P).$$

Knowing which, that is, what entity fits a description, is

$$Knowref_A(P(x)) \Leftrightarrow (\exists y)((\forall z)P(z) \Leftrightarrow y=z) \wedge B_A((\forall z)P(z) \Leftrightarrow y=z).$$

In other words, there is a unique value for  $x$  making Prop true, and A believes that

this value uniquely satisfies Prop. It is possible to want (W) either an action or a proposition; agents believe that they Want what they Want, and Wanting is required to be distributive over conjunction.

Action types have a name, a set of constrained parameters, and formulas labelled Effects, Body, and Preconditions. The Body is a list of goal states rather than subactions. A Plan to transform one world into another is a sequence of actions such that each action's preconditions hold in the preceding world, and the action transforms that world into the current one. Agents believe that actions achieve their effects and require their preconditions. Any action that occurs was intended (W) by the agent.

Agents model each other's plan construction and recognition processes by chains of plausible (non-deductive) inferences. There are four plan construction rules, and five corresponding recognition rules:

$W_A(P) \rightarrow_e W_A \text{Knowif}_A(P)$   
if an agent wants a proposition, she may want to KNOWIF it holds [KNOWIF rule]

$W_A(Y) \rightarrow_e W_A(X)$ , X a precondition of Y  
if an agent wants an action, she may want its preconditions [action-precondition]

$W_A(X) \rightarrow_e W_A(Y)$ , X an effect of Y  
if an agent wants a proposition, she may want an action having this effect  
[effect-action]

$W_A(Y) \rightarrow_e W_A(X)$ , X a step of Y  
if an agent wants an action, she may want its body [action-body]

$B_S W_A \text{Knowif}_A(P) \rightarrow_i B_S W_A(P)$   
if the system believes the agent wants to Knowif P, it infers she may want it to be true [know-positive]



$B_S W_A \text{Knowif}_A(P) \rightarrow B_S W_A(\neg P)$

if the system believes the agent wants to Knowif P, it infers she may want it to be false [know-negative]

$B_S W_A(X) \rightarrow B_S W_A(Y)$ , X a precondition of Y

if the system believes the agent wants a proposition, it infers she may want an action with this precondition [precondition-action]

$B_S W_A(Y) \rightarrow B_S W_A(X)$ , X an effect of Y

if the system believes the agent wants an action, it infers she may want its effects [action-effect]

$B_S W_A(X) \rightarrow B_S W_A(Y)$ , X a step of Y

if the system believes the agent wants the body of an action, it infers she may want the action [body-action]

A special case of the precondition rule is, if the system believes the agent wants another agent to want an act, she may herself want that act [want-rule].

Speech act theory requires that speakers intend these intentions themselves to be recognized, so Perrault and Allen add schemas embedding each side of a rule. Nested plan construction rules embed each side of each rule above in  $W_{\text{Speaker}}(\dots)$ . Nested recognition rules embed each side in  $B_{\text{Hearer}}(W_{\text{Speaker}}(\dots))$ . An agent can even plan for another agent to construct a plan, and intend for the other agent to recognize this. This is done by embedding the original plan construction rules twice. The corresponding inference space is explored by heuristic search until an action description is identified which fits the observations, context, and expectations.

The heuristics used in the search are again based on the structure of the actions. They favor actions with true preconditions, those with false effects, those whose effects are intended, and those which the agent is actually able to perform. The following example shows an inference chain which is the most favored by the

heuristics, but does not show the heuristics themselves. We simply note in advance that each heuristic supports the plan being considered at each step.

As an example, let us consider the Spanish example that was discussed extensively in Ch. 2. The pure inference method requires two speech act definitions: an S-REQUEST or surface request, and an ordinary Request. Surface acts are associated directly with the mood of the sentence, and since questions are treated as Requests to Inform, S-REQUESTs comprise imperative and question sentences. They simply have the effect that the hearer believes the speaker wants the hearer to perform an action. A genuine Request has the precondition that the speaker want the hearer to perform the act, and the effect that the hearer wants to perform the act. The body of a Request matches the effect of an S-REQUEST, so that an S-REQUEST is one way of actually Requesting.

Suppose that Mrs. de Prado (P) and Suzanne (S) mutually believe (MB) that S can speak Spanish. Mrs. de Prado says "Can you speak Spanish, please?" The Allen method does not take advantage of the cue "please", but begins with the intended literal question. SBPW should be read as "S believes P wants". The initial belief triggered by the utterance is

SBPW(S-REQUEST(P,S,

INFORMIF(S, P, ABLE(S, SPEAK-LANGUAGE(S, Spanish))))

Now, since an effect of a REQUEST is that the hearer perform the REQUESTed action, S can use the action to effect rule, to conclude that P wants it to be well-known that P wants to be informed.

SBPW(MB(S, P, PW(

INFORMIF(S, P, ABLE(S, SPEAK-LANGUAGE(S, Spanish))))))

The effect of an INFORM is that the hearer KNOW the proposition, so the mutual-belief rule linking actions to their effects yields

SBPW(MB(S, P, PW(

KNOWIF(S, P, ABLE(S, SPEAK-LANGUAGE(S, Spanish))))))

If you want to know something, it might be because you want it to be true.

[know-positive rule.]

SBPW(MB(S, P, PW(ABLE(S, SPEAK-LANGUAGE(S, Spanish))))))

If you want a precondition of an action, you might want the action. [precondition-action rule.]

SBPW(MB(S, P, PW(SPEAK-LANGUAGE(S, Spanish))))

This is the body of a request, in Allen's scheme, so Suzanne can reason from the body to the action's identity as a request.

SBPW(REQUEST(P, S, SPEAK-LANGUAGE(S, Spanish)))

This chain of reasoning is favored by the set of recognition heuristics. The yes-no question interpretation would arise by reasoning from body to action after conclusion 2. However, it is discounted by the heuristics because its effects already hold in this context. Other possible interpretations also conflict with the context or make less use of mutual belief. The same chain of reasoning applies any time an action precondition is queried. The crucial link could equally well be based on any other plan reasoning or causal rule, however.

It is important to note the distinction between responses that a hearer makes in order to be helpful, and responses that the hearer gives after recognizing that this is the desire that the speaker intended to communicate. If the speaker says, "the table is dirty", you might infer that the speaker wanted you to clean it, and for you to recognize that. This is recognizing a request. If the speaker only meant to warn you not to set anything in the mess, you can still clean the table out of helpfulness (or out of your own interests.) But this does not make the warning a request. Individual cases may have elements of both, of course, but different paths of reasoning are involved.

#### 4.1.4. Discussion

We now examine how such a plan-based approach to speech act interpretation plays the four roles mentioned at the beginning of this chapter. These were

- 1) eliminating speech act interpretations proposed by the linguistic mechanism, if they contradict known intentions and beliefs of the agent.
- 2) elaborating and making inferences based on the remaining interpretations, allowing for non-conventional speech act interpretations.
- 3) proposing interpretations of its own, when there is enough context information to guess what the speaker might do next.
- 4) providing a competence theory motivating many of the conventions we have described.

This plan-based approach is very powerful and very general, and is based on mechanisms needed by agents whether or not they communicate. The heuristics direct the search toward interpretations which are plausible in this context and away from those which are not. They thereby provide a partial ordering on

interpretations, based on certain components of the context. Just as they chose the Request interpretation over a yes-no interpretation in our Spanish example, they can order proposed interpretations in any context. As always, there are two important components to this process. One real strength of plan reasoning is in the knowledge representation: plan definitions enumerate facts which are relevant to speech act plausibility. Not only are the conditions on the speech acts relevant, but the conditions on the acts they describe are relevant also. Any formalism that is adequate for planning incorporates the most relevant information about possible actions, and hence provides an index into contextual factors. The second strength of plan reasoning is the interpretation component: any planning system has the information, but how this information is used is also crucial. Heuristic search was subsequently used by [Sidner 81]. [McCafferty 86] and the current work emphasize using the heuristics to add information to the system. The current work further emphasizes the screening process over the search process.

Plan reasoning is also very useful for explaining non-conventional speech acts. It is precisely these that require the full generality of the mechanism. Suppose you are in a car, by the only open window, and another passenger says "It's cold in here." Assume it's well known that a cold car causes the agent to be cold, that it is bad for agents to be cold, and that an open window can make the car cold.

The plan reasoning is as follows.

SBAW(S-INFORM(A, S, Cold(space1)))

SBAW(MB(S, A, AW(S KNOW Cold(space1)))) (action-effect)

SBAW(MB(S, A, AW(S KNOW Cold(A)))) (causal)

SBAW(MB(S, A, AW(S W not(Cold(A)))) (undesireability)

SBAW(MB(S, A, AW(S W not(Open(window1)))) (planning by causal)

SBAW(MB(S, A, AW(S W Close(S,window1)))) (planning by  
effect-action)

SBAW(MB(S, A, AW(Close(S,window1))) (want-action)

SBAW(Request (A, S, Close(S,window1))) (body-action)

In other words, you are to know that the car is cold, so the speaker is cold, but that's bad. You can plan to fix it by closing the window, so the speaker wants you to *want* to do it, so the speaker wants you to do it and is therefore requesting that you close the window. The hearer need never have heard this request before, nor even one requiring similar reasoning. All we need is this domain plus the general principles.

The sense in which this utterance is specifically a request to close the window depends crucially on the simplicity of the planning step. It is possible in this limited car environment to mutually believe that the problem is the open window, and not the air conditioning or the choice of locality. We will be concerned later with simplicity. For now, it is important to note mainly that the speaker can count on the hearer to perform such plan construction. Any agent that can reason about plans and other agents can understand a great variety of novel speech acts. The same reasoning that provides new interpretations may simply elaborate a more direct act.

Expectations about the speaker's plans have not been an explicit part of speech act interpretation models. However, it would be an obvious extension of existing work. The Allen system took advantage of the limited domain to keep the set of possible plans small. Subsequently Kautz showed [Kautz 87] what the theoretical limits are on plan recognition based on a hierarchical plan library. The plan library is organized into a taxonomy by an abstraction relation, and each action is connected to its steps by a decomposition relation. The input is a series of observed steps, and a search of the hierarchy yields a list of the possible top level plans in progress as well as a list of possible next actions. If there were a plan for question-answer pairs, for example, observing a question would suggest a question-answer pair in progress. The system could then try to interpret the next input as an answer. Plan tracking is an important part of the discourse systems of [Litman 85] and [Grosz 86a]. Tracking of domain plans only is pursued in detail by [Carberry 87].

Although plan reasoning ignores the conventional aspect of surface form, it is one important motivator of form. "Please" itself is the residue in American English of a happy condition explicit in French requests. "If you please", meaning "if it pleases you" (see Ch. 3), is an alternative form of the precondition that an agent wants an action. In this case it is the precondition on the act being requested, and so the plan-based approach motivates our lexical convention. Gordon and Lakoff's generalizations about querying vs. asserting felicity conditions of actions can also be motivated on plan reasoning grounds. Felicity conditions are approximately the

preconditions of the act. They can be queried when the hearer is assumed to be the authority for that fact, and asserted when the speaker is. For example, the Hebrew "You want to make me some dinner." request asserts a precondition of the requested action, and as such fits the plan reasoning approach readily. It strikes Americans as presuming, to inform people of their own wants; this argument too can be stated in plan reasoning terms.

We see that a linguistic account of conventional speech acts leaves important work to be done by plan reasoning. We will complete our survey of plan-based theories before discussing the sort of plan reasoning we have in mind.

## 4.2. Related Work

### 4.2.1. Perrault

One of the difficulties of speech act theory is the morass of nested beliefs and intentions which are necessary to differentiate communication from causality, and to explain complications like irony and lying. The original insight about beliefs and communication is Grice's [Grice 57]. Communication depends crucially on a reflexive intention. The speaker must intend to produce some effect in the audience, *by means of the recognition of this intention*. Agents may not communicate when making statements to test a microphone; they intend to produce an acoustic effect by physical means. In communication, physical means are necessary but not sufficient; the hearer must also believe that the speaker wants the information transfer. Further, it is not enough to suspect that you were meant to



overhear a remark, either; the speaker must *overtly* intend you to believe. Then if you remain skeptical about the information, you still recognize the attempted communicative action.

Austin [Austin 62] distinguished three kinds of act. Locutionary acts are the uttering of words and sentences. Illocutionary acts are done by performing some locutionary act in a particular context with particular intentions. Perlocutionary acts are roughly the consequences of the previous two; getting someone to believe something, as opposed to telling them. Perlocutionary acts need not be intentional. Work on speech acts is concerned primarily with illocutionary acts. And in order to model the communicative intentions in illocutionary acts, Perrault and Allen resorted to three levels of embedding and a claim like this: for S to perform an illocutionary act IA,  $W_S B_H W_S B_H W_S (E)$ , where E are the effects of IA. The effects of IA for a request would be  $W_H (Do(H, A))$ . The nested beliefs and intentions grow cumbersome. Perrault's default theory of speech acts [Perrault 87] provides an elegant approach to this problem.

Perrault rightly notes that speech act effects are highly dependent on the beliefs that agents have already. Thus, these effects are best regarded as defaults only, which can be defeated in the presence of conflicting information. He models how agents may revise their beliefs after a speech act, using Reiter's default logic. ([Reiter 78] provides a logic in which inference rules are defeated --rendered inapplicable-- by the failure of an associated applicability condition. This results in a model theory with different extensions for different inference orderings.)

Agents in Perrault's logic are modelled as follows. They remember their beliefs over time and continue to hold them. If they observe an action they believe it was done. In addition to these axioms there are two default rules. One states that an agent can acquire a belief if the agent believes another agent holds it. The second rule is specific to declarative sentences, and says that uttering a declarative sentence implies that the speaker believes its contents. Thus if S says to H that the sky is blue, H reasons that the sky is blue.

$Do_{H,0}Obs(S)$	H observes S at time 0
$B_{H,1}Do_{S,0}p$	H noted S's declarative utterance of p
$B_{H,1}B_{S,0}p$	H infers S believed p
$B_{H,1}B_{S,1}B_{S,0}p$	H infers S remembers believing p
$B_{H,1}B_{S,1}p$	H infers S continues to believe p
$B_{H,1}p$	H decides to believe p too.

So now H believes that the sky is blue. If S observed S and H at that same time, S can reconstruct H's default reasoning, and H can reconstruct this reasoning of S's, and so on ad infinitum.

Thus from a very simple formulation, it is now possible to infer many of the complicated beliefs that we need for a successful account of communication. It can model lying, by adding that the speaker simply doesn't believe the statement, and cannot be convinced by the statement or by the hearer's new beliefs, because the default rules will be defeated for the speaker only. Perrault's logic is a very concise statement of the mechanism, because it leaves the nestings to be constructed by the derivation process. And to be precise, it is then necessary to

define speech acts in terms of infinite series (we need belief integrals.) This does indeed incorporate dependence of beliefs on the previous mental state of the agents, but has yet to handle intentions.

#### 4.2.2. Cohen&Levesque

Cohen and Levesque have pursued a similar line of work, although they do not incorporate a theory of default reasoning. They developed a formal notion of commitment, allowing them to express goals which persist until they are either satisfied or obviated [Cohen 86]. Goals may be formulated which are conditional on arbitrary propositions, allowing them to be dropped if the situation changes. For instance, if it rains one might decide not to water the garden after all. This allows Cohen and Levesque to formulate speech acts in such a way that the agent is committed only to being understood, not to any particular speech act. Also, they can express the fact that the hearer of a request may abandon the requested action, if the hearer realizes the speaker no longer wishes it to be done. Cohen and Levesque regard it as an advantage of their approach that the speech act classes themselves are epiphenomena [Cohen 88].

These developments in speech act theory and knowledge representation are substantive and foundational. One would like to know how they can be extended to accomodate much richer linguistic information, for speech act recognition. (Our method makes use of explicit speech act representations, and therefore could not be integrated directly.) One would also like to be sure that the methods scale to a full

range of speech acts: what would a greeting look like, for instance, and would adding it obscure or invalidate the mechanisms? To date, Perrault handles only Inform and Cohen and Levesque only Request. The hope would be that these approaches can provide tools and clean speech act definitions to the next generation of speech act recognizers. Recognizers themselves, we claim, will require more explicit and detailed information and less inference.

#### 4.2.3. Kautz

Kautz's *Formal Theory of Plan Recognition* [Kautz 87] includes a speech act example. Kautz defines an abstraction hierarchy based on a reified logic of events. He then provides a method based on circumscription, for identifying what plans may be in progress based on observations of primitive actions. This method has a model theory in which the set of unrelated observations is minimized. For speech acts, the primitive actions were surface speech acts corresponding to sentence mood. Then these were listed as decompositions of various illocutionary acts, which were in turn part of other plans involving language. The algorithm takes a Surface Request, for example, and sees that it may decompose a Question or an Indirect Request. These in turn may be part of a plan to get information or a plan to have the hearer do something. The algorithm then checks constraints, rejecting interpretations whose associated plans are implausible. The method's attraction is its clean semantics. It also makes use of information about plans which may be in progress. However, extension of the theory to handle linguistic features more

appropriately would create very high branching factors at the leaves of the hierarchy, if indeed the semantics can be sustained.

### 4.3. Short Inferential Distance

As do the previous approaches, we emphasize the hearer's model of the speaker's intentions. Even more so than Allen and Perrault, we treat these intentions as conclusions to be drawn from the utterance rather than facts known beforehand. Furthermore, we emphasize short inferential distance in the conventional cases, relying on our notion of plan-based conversational implicature. Rather than using extensive search to determine what is proved, we base decisions about speech act interpretations on the small, finite list of beliefs associated directly with their definitions. One could compare this roughly to some fixed number of breadth-first plies of Allen & Perrault's rules, or to the database checking that the Kautz algorithm would do for speech acts if they were treated as ends in themselves.

The plan reasoning component of our approach assumes that there are several dozen standard illocutionary acts like Request and Greet. These are represented as as primitive actions in a plan hierarchy with abstraction and decomposition relations. Plan reasoning takes an illocutionary act as input, and returns a set of inferences. The inferences resemble Allen and Perrault's search heuristics, or Kautz's constraint checking. For an illocutionary act we will attempt to prove preconditions, constraints, and other related propositions. We attempt to prove these things not with respect to the absolute truth, but with respect to the hearer's

model of the speaker's beliefs. This move follows directly from a concern with speaker meaning [Grice 71]; the speaker may very well be trying to inform us of a notion we already hold, yet we still recognize the intent.

When we attempt to prove one of these propositions, there are three possible results: true, false, and unknown. If it is true, this is evidence for our speech act interpretation. If false, it is evidence against. But if the hearer does not know what the speaker believes ( $\neg \text{Knowif}(H, SB...)$ , not  $HB \neg \text{Knowif}(S....)$ ), the action interpretation is itself evidence for the belief. The fact that our knowledge is not complete is one motivation for regarding these beliefs as new information. A second motivation is a strong resemblance between these inferences and the plan-based subset of conversational implicatures.

In the rest of this section we will see how the method serves to test speech act interpretations in context, serving the first purpose mentioned for plan reasoning in this chapter. In later chapters we will see that this plan reasoning process can be used as a filter, weeding out inconsistent interpretations and identifying ones for which there is already evidence.

#### 4.3.1. Plan-Based Conversational Implicature

The problem of conversational implicature, we recall from Chapter 1, concerns conclusions drawn from an utterance, which are not justified by classical logic because they are based on defeasible assumptions about rational behavior. Recall Grice's example [Grice 75]:

A is standing by an obviously immobilized car and is approached by B;  
the following exchange takes place:

- (70) A: I am out of petrol.  
B: There is a garage round the corner.

(Gloss: B would be infringing on the maxim 'Be relevant' unless he thinks, or thinks it possible, that the garage is open, and has petrol to sell; so he implicates that the garage is, or at least may be open, etc.)

B has communicated much more than the location of a station. If B knew it to be closed, B's reply would be misleading. Having recognized that A's goal was to get some gas for stranded car, B took into account the preconditions of buying gas. Plan reasoning provides the links that augment this utterance to the point of relevance.

These conclusions are clearly based on the participants' goals; if A had a flat, B's implication that the garage has gas would be displaced by having the appropriate tools and so on. But these are just conditions on the corresponding domain plan:

Plan-based conversational implicatures include those beliefs and intentions that contribute to having a plan. Specifically, the speaker must be willing to believe

**header:** Buy-Gas(agent, seller, loc, time, gas)  
**preconds:** OWN(agent, price(gas))  
**constraints:** OWN(seller, gas) AT(seller, loc, time) AT(gas, loc, time)  
**decomp:** Goto(agent, loc, time)  
           Give(agent, seller, price(gas), time)  
           Give(seller, agent, gas, time)  
**effects:** OWN(seller, price(gas)) OWN(agent, gas)  
           -OWN(agent, price(gas)) -OWN(seller, gas).

#### Gas Plan

that the constraints hold, that the preconditions are satisfied or can be, and that the effects will hold in the end. The speaker believes the agent intentionally does the steps, wants the goals, wants the action, and so on. Later we will need to be careful about how we specify the exact agents and times, but for now we will just state the main ideas informally. To plan at time  $t_1$ , for an action at time  $t_2$ , the agent must believe

- that the plan's constraints will hold at  $t_2$ ,
- that the plan's preconditions can be achieved by  $t_2$  and that the agent intends to achieve them,
- that each of the actions in the decomposition is performable at  $t_2$
- that each has some useful role in the plan, and that the agent actually intends to do them at  $t_2$ ,
- and that the effects of the plan will hold after  $t_2$ , which would not be the case were the plan not executed.

These beliefs of the agent are similar to those given by Pollack for purposes of plan recognition in question-answering, where the speaker may have a faulty plan. Pollack makes use of Goldman's generation and enablement relations [Goldman 70] rather than the STRIPS model of plans. Since the explanatory power of her ideas depends generally upon having a good plan representation, our mechanisms too should be adequate to support this kind of reasoning about erroneous plans. For the moment we would like to show that it yields an interesting class of conversational implicatures.

For the gas station example, A infers that B believes

- There are a seller, a location, a time and some gas.  
(The variables in the plan have reasonable bindings.)



- The agent has some money or can plan how to get some.  
(The plan preconditions hold or can be achieved by plans.)
- The seller owns the gas, and both are at the gas station location at the time.  
(Plan constraints will hold at the time of plan execution.)
- The agent need only go there, hand over the money, and receive the gas.  
(Plan decomposition is appropriate and workable.)
- Then the seller will own the money and the agent, the gas.  
(The effects of the plan will hold after its execution.)
- There isn't anything likely to interfere with this plan.  
(The effects of the plan will hold after its execution.)

Thus a small set of parameterized inference rules, when applied once to the plan, yields the specific conclusions that were indicated by Grice. We refer to conversational implicatures like this example as plan-based conversational implicatures [Hinkelman 87].

Plan-based implicatures meet all of the criteria for conversational implicature. An implicature is neither a truth condition nor an entailment of an utterance, in the classical sense. Rather, it depends on assumptions about cooperative agents and their ability to act "rationally". This is clearly the case with plan-based implicatures; it is only when we assume a logic of action that we can make these inferences. And this logic of action is not a formal property of the universe but a description of human behavior. Plan-based implicatures are *cancellable*, that is, it is possible to assert a sentence but deny its implicatures, without logical contradiction. For instance, one could coherently say, "There's a gas station around the corner, but it's probably closed." They are *detachable* from the utterance; any utterance with the same classical semantics should have the same conversational implicatures in the same context. Grice's test for detachability is paraphrase. A conclusion that hinges on a particular word in the utterance is not conversational but rather a *conventional* implicature. (There are some arguments

about detachability as an implicature criterion: it is aimed at isolating lexical connotations and doesn't hold up under phenomena like topicalization, for instance.) "Go one block north and half a block west, and you'll see a gas station" does not differ from the original in its implicatures. Conversational implicatures are regarded as being intentionally communicated. Plan-based implicatures by no means account for every conversational implicature, but they do yield a major class of implicatures.

#### 4.3.2. Implicature and Speech Acts

We have just seen that there are certain inferences which have a strong but defeasible connection to an utterance in context, and that some of these conversational implicatures are plan-based. We saw that they are very closely related to speech act interpretation, but we have not yet elucidated the exact nature of this relationship. An approximate answer is this: while speech act interpretations themselves constitute defeasible inferences from an utterance, and in that sense may be regarded as implicated, we will reserve the term *implicature* for inferences which are derivative of a particular speech act interpretation. However, such inferences are defeasible with respect to the utterance but not with respect to the speech act, so that they also act to filter out implausible speech act interpretations.

For an agent to perform a speech act sincerely, the agent must hold the appropriate beliefs about both this discourse plan and about any domain plan. These beliefs

are the plan-based implicatures described above. Consider the Spanish example once again, recalling that the utterance occurs in a context where it is mutually known that Suzanne speaks Spanish. This was the plan for the Ask interpretation, which we will call A1:

```
(ASK-ACT AGENT Mrs. de Prado
  HEARER Suzanne
  PROP (ABLE-STATE AGENT Suzanne
    ACTION (USE-LANGUAGE AGENT Suzanne
      OBJECT ls1)))
```

Some implicatures for the ASK act are shown below.

```
P2 = (ABLE-STATE AGENT Suzanne
      ACTION (USE AGENT Suzanne
        OBJECT ls1))
```

**from Effects**

Want(Mrs. de Prado, Knowif(Mrs. de Prado, P2))

**from Standard Preconditions**

Believe(Mrs. de Prado, ~ Knowif(Mrs. de Prado, P2))

Believe(Mrs. de Prado, Cando(Mrs. de Prado, A1))

**from Precondition**

Believe(Mrs. de Prado, Knowif(Suzanne, P2))

**from Action**

Intend(Mrs. de Prado, Utter(Mrs. de Prado, "Can you speak Spanish?"))

**Implicatures for A1 = ASK(Mrs. de Prado, Suzanne, P2)**

Since in context it was well known that Suzanne speaks Spanish, the ASK act's second implicature under Standard Preconditions is implausible. So the ASK

interpretation is eliminated. (Normally we would need to compute the implicatures for the embedded action, but the contradiction allows us to give up.) The set of implicatures for this speech act interpretation has served to identify the contextual conditions under which this speech act interpretation would be plausible, and has thereby pinpointed the implausibility of this particular interpretation in this particular context. They have filtered out this interpretation.

Conversational implicature relies on extended plan reasoning for its own competence theory. But what it does for speech act recognition is to provide the link to relevant context, at reasonable cost. Thus our conversational implicature mechanism provides sufficient plan reasoning capability to constrain speech act interpretation greatly.

We use pragmatic inferences such as plan-based implicature and presupposition as a restricted variety of plan inference that acts to filter the speech act interpretations, reducing ambiguity as well as yielding the implicatures. Extended reasoning about plans, as exemplified by Perrault & Allen, will still be required for novel speech acts and for a competence theory, but need not be invoked in the majority of cases.

In the next chapter we will introduce the machinery more formally, with a full speech act hierarchy and definitions. Then we will be in a position to look at some extended examples.

## 5. Plan Reasoning II

In the previous chapter, we showed that general reasoning about plans can contribute to speech act recognition in several different ways. We showed how it has been used to derive interpretations of novel acts. We then focussed on its role in screening interpretations, and showed how a shallow search through a limited inference space serves not only to screen out implausible interpretations, but yields useful inferences. This chapter specifies plan reasoning more precisely. It provides an inheritance hierarchy of speech act definitions, and uses the definitions as a basis for examples of the reasoning in action.

### 5.1. Knowledge Representation Issues

Plan representation for natural language processing need not be done in the STRIPS tradition. [Pollack 86], for example, relies on Goldman's generation and enablement relations [Goldman 70], for reasoning about misconceptions in speakers' plans. We contend that any representation which is adequate for planning will, with some representation of belief and intention, suffice as a basis for speech act recognition. However the STRIPS representation has been well studied, and was used for many results that we draw on, so we will continue to use it here. [Tenenbergs 89] provides a formal account incorporating inheritance abstraction into STRIPS-style planning systems, with well-defined semantics.

### 5.1.1. The Logic

While extended reasoning requires the definitions of Chapter Four, our plan-based implicature computation is somewhat simpler. The belief operator  $B(A, P)$  is just the one defined in Chapter Four. But since we make no use of the notion of objective truth, we omit the knowledge operator  $K(A, P)$ . We weaken the *Knowif* and *Knowref* operators correspondingly, to indicate only that the agent holds some belief about the subject. *Knowif* is used to represent our belief that another agent has an answer to a yes/no question, when we ourselves do not know which answer that is:  $Knowif(A, P) \Leftrightarrow B(A, P) \vee B(A, \neg P)$ . The analogue for wh-questions, knowing which entity fits a description, is  $Knowref(A, P(x)) \Leftrightarrow (\exists y) B(A, (\forall z) P(z) \Leftrightarrow y=z)$ . In other words, for some value, A believes that this value uniquely satisfies the description P. The intention operator  $W(A, X)$  is that of Chapter Four, with the additional requirement that agents do not want both a state and its negation:  $W(A, P) \Leftrightarrow W(A, \neg P)$ .

$$B_A(P) \rightarrow B_A(B_A(P))$$

$$B_A(P) \wedge B_A(Q) \rightarrow B_A(P \wedge Q)$$

$$B_A(P) \vee B_A(Q) \rightarrow B_A(P \vee Q)$$

$$B_A(\neg P) \rightarrow \neg B_A(P)$$

$$(\exists x) B_A(P(x)) \rightarrow B_A((\exists x) P)$$

$$(B_A(P \rightarrow Q) \wedge B_A(P)) \rightarrow B_A(Q)$$

It is also closed under Modus Ponens and the axioms.

$$\text{Knowif}(A, P) \Leftrightarrow B(A, P) \vee B(A, \neg P).$$

$$\text{Knowref}(A, P(x)) \Leftrightarrow (\exists y) B(A, (\forall z) P(z) \Leftrightarrow y = z).$$

$$W(A, P) \rightarrow W(A, \neg P).$$

The definition of an action is more complicated. The type definition of an action includes a name, a set of constrained parameters, and formulas labelled Effects, Body, Preconditions, and Constraints. The type definition of an action has variable parameters, while an instance of an action has only constant values. An action Body consists of a list of actions, a list of states, or nothing. The list of actions is a set of steps which achieve the parent action, and their temporal ordering if any must be specified by constraints attached to the parent plan. The Decomposition relation holds between any step and the parent. The list of states specifies a set of goals whose achievement constitutes a performance of the parent act; this is included for compatibility with the method of Perrault & Allen. If no Body is given, the action is realized by processes about which the system does not currently reason, for example, the linguistic process for speech act generation.

Preconditions and Constraints are propositions which must hold in the current world state in order for the action to make the effect propositions true. Constraints are propositions which are normally out of the agent's control, while agents may plan to achieve preconditions. The Choosable predicate of Pelavin [Pelavin 86]

can be used in a given context to make this distinction: it states that that for any possible world there is some series of actions which the agent can perform to bring about the precondition. (The precondition may even be inevitable under this definition.) Agents believe that actions achieve their effects and require their preconditions and constraints. The Abstraction relation holds between any pair of actions such that if the second occurs, with its preconditions, effects, and so on, it follows that the first has occurred, with its preconditions, effects, and so on. This relation allows us to build a hierarchy of actions which can be used as a basis for reasoning even when not all information about an act is known at this point.

### 5.1.2. Notation for Actions

In the linguistic chapters, we denoted actions essentially by their headers. We used a representation of categories with slots and fillers, in which the category was the action type, and the slots were essentially typed variables, to be filled with constants. Now we will condense this notation, and indicate the types of variables and constants explicitly by separating the type name from the identifier with a colon. For now, all actions are action types, with variables. Only in particular contexts will we discuss action instances with constants for arguments. We also represent other action components explicitly. Occasionally we may write *Preconditions(A)*, *Effects(A)*, and so on to denote the set of propositions or objects having that label.

**Plant(S:Human,T:Seed)**

**Preconditions: Has(S, T), At(S, G:Garden)**



Constraints:  $\neg \text{Dead}(T)$   
 Body:  $\text{Dig}(S, H:\text{Hole}), \text{Put}(S, T, H), \text{Cover}(S, T, D:\text{Dirt})$   
 Effects:  $\text{Sprout}(T)$

This is the action description for planting a seed. The agent must plan to have a seed and to get to the garden. The agent carries out the steps of digging the hole, placing the seed in the hole, and covering it. Under the constraint that the seed is actually alive to begin with, (and other qualifications!) it sprouts.

For convenience we omit from the header the less important arguments to the propositions; this is for presentation purposes only as the knowledge representation so far has no provision for local variables. There are some variables which we will use habitually: S stands for 'speaker' and has type Human, H stands for 'hearer' and has type Human, P stands for 'proposition' and has type State, and A for 'action' and has type Voluntary-Action. The proposition  $\text{Do}(H, A)$  is used essentially for type conversion, it is true if action A occurs with agent H, and thus allows A to be included as a qualification on the definition of another action. For instance, A might be the action requested by a Request, and is therefore Done as an effect of the Request. The action  $\text{Achieve}(H, P)$  is an action with H as the agent and P as an effect. We thus have

$$\text{Do}(H, \text{Achieve}(H, P)) \Leftrightarrow P$$

$$\text{Achieve}(H, \text{Do}(H, A)) \Leftrightarrow A$$

We do not at this time provide a calculus of higher-order beliefs and intentions, although one would be desirable.

### 5.1.3. Plan-Based Implicatures

We have previously made very general comments about the defeasible inferences that planning agents can make based on actions. Here we specify a basic set of inferences which will go a long way toward the elimination of implausible speech act interpretations. In our definition of actions, we noted that agents believe that actions require their preconditions and constraints, and yield their effects. This can be restated as follows:

$Do(S, A) \rightarrow B(S, X)$ ,  $X$  a precondition of  $A$   
 $Do(S, A) \rightarrow B(S, X)$ ,  $X$  a constraint on  $A$   
 $Do(S, A) \rightarrow W(S, X)$ ,  $X$  an effect of  $A$

We can add one simple fact about agents, namely that they do actions because their effects do not already hold:

$Do(S, A) \rightarrow B(S, X)$ ,  $X$  an effect of  $A$

A more temporally sophisticated version of this rule could be based on Pelavin's [Pelavin 86] Inevitable predicate, and it would state that the agent performs an action because one of its effects is not inevitable. This would allow expression of plans for the maintenance of some condition which does hold at the time of maintenance. However, we use a simple view of time throughout, in which the inferences are computed after the beginning of action execution but before the effects have been secured. This view is particularly appropriate for speech act understanding, in which these inferences must be computed in order for the effects to be achieved. The desired effects may even fail to result.

These inference rules are essentially action recognition rules. When we model an agent who is reasoning about a second agent, we must embed them into the agent's belief space:

$B(H.Do(S.A)) \rightarrow B(H.B(S.X))$ , X a precondition of A  
 $B(H.Do(S.A)) \rightarrow B(H.B(S.X))$ , X a constraint on A  
 $B(H.Do(S.A)) \rightarrow B(H.W(S.X))$ , X an effect of A  
 $B(H.Do(S.A)) \rightarrow B(H.B(S.X))$ , X an effect of A

The results of these inference rules will be our implicatures. We do not require that the agent explicitly plan to communicate them, nor that this intention itself be recognized. (This is a departure from Grice.) It is simply a part of the communication process that these are derived, and that the speaker relies on the hearer to make such a computation. The speaker need not enumerate these things explicitly. Grice's qualification "as far as the speaker knows" is taken to be adequately captured by our explicit representation of the speaker's belief. If it proves to be too strong a statement about the speaker's beliefs, we can fall back on consistency:  $\neg B(S, \neg X)$ .

This set of rules is a simple one, and does not include many aspects of plan structure that it could. For instance, it may be possible to eliminate some speech act interpretations based on what we know about the act's uses in the decomposition hierarchy, or on whether it is possible to find values for all the variables. The agent should also believe that the preconditions were achieved rather than inevitable, and the reverse for constraints. There are also the causal connections represented explicitly by Pollack: an agent must intend any steps as

*part of* the action, not merely as part of some other action. There will be some such set of details for any particular knowledge representation chosen. The general notion of a set of beliefs about a plan remains constant.

## 5.2. A Speech Act Taxonomy

We next present an inheritance hierarchy of speech acts, with several goals in mind. First, it must be adequate to account for a wide range of ordinary dialogue, including any examples we wish to discuss. Second, the definitions must capture classes of actions in a way adequate for planning as well as attributing these actions to other agents. Third, the categories should be intuitive and illustrative of the type of knowledge representation needed for speech act recognition.

This figure shows the abstraction relations at the top levels of a hierarchy of speech acts. For simplicity we show just their types and parameters, and discuss their full definitions below. The class of speech acts is a subtype of voluntary actions, and it subdivides into five main categories taken from [Searle 79]. Representative acts are those in which the speaker indicates some belief about the world's state, regardless of the degree or accuracy of the belief. Informing, speculating, and boasting fall into this category. Directive acts are attempts to get the hearer to do something; requests and commands are the paradigm examples. Commissive acts are those in which the speaker is bound to bring about a state of the world, and promises are prototypical commissives. Expressive acts, such as condolences, are nominally expressions of attitude about some state of affairs, and not in general

attempts to achieve something or describe the world. Searle contrasts them with declarative acts, which comprise explicit performative acts like "You're fired!". He says that declarative acts do create the state of affairs that they mention. The distinction is more useful than airtight. The sixth class shown here is not one of Searle's. He treated questions as Requests to Inform. Although their logical structures are closely related, their linguistic differences are great enough to make the distinction worthwhile. Those differences were discussed in Chapter Three.

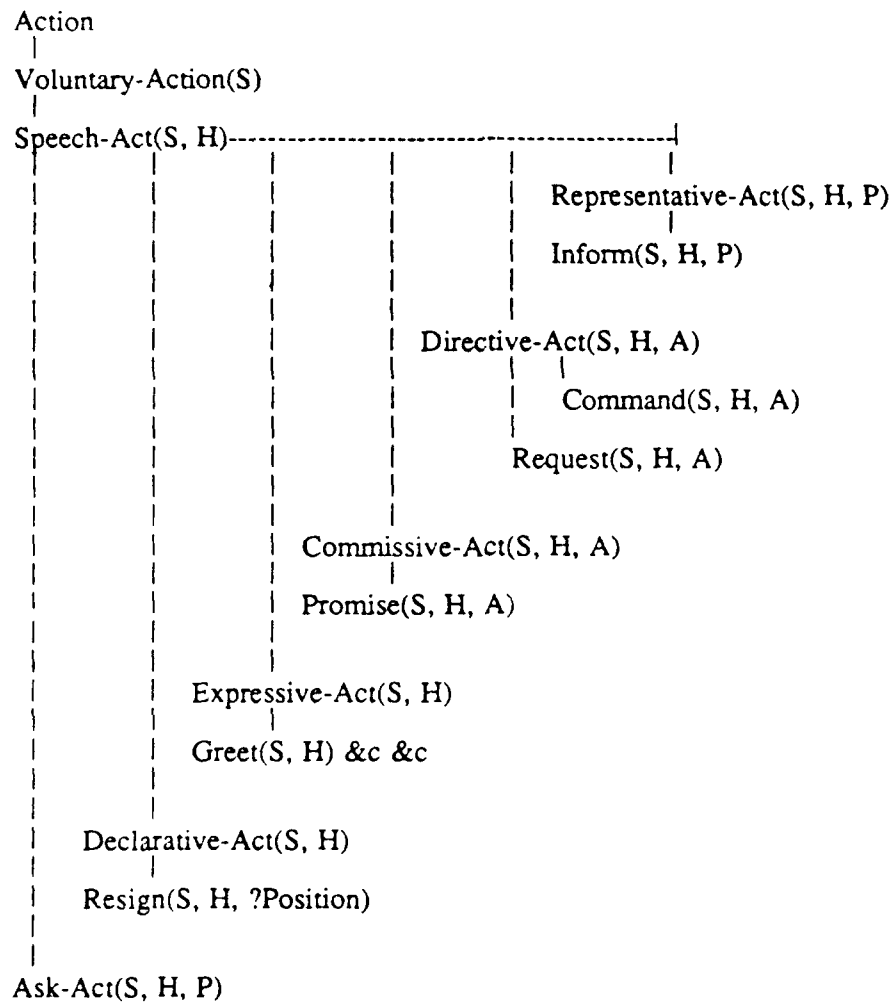
We now provide definitions for the actions. The only generalizations that can be made about all actions are that the agent must be animate, and that the agent must be able to perform the action. This is an abstraction of the specific capabilities that the agent would need to perform the act, and the specific conditions which must hold for the action to be successful. Any executable action contains these specifics in its constraints. A voluntary action is simply an action which is done intentionally. The class of speech acts has general observation conditions: that the speaker and hearer are actually paying attention to each other. The observation conditions of Perrault are here specialized. We simply assume that the agents use the same language and have the appropriate sensory abilities for the communication medium. They are necessary to any real application because they are not valid in general, but they would serve here only to clutter our examples.

**Action(S:Animate)**

**Preconditions:**

**Constraints:** Able(S, Self:Action)

**Voluntary-Action(S:Agent)**



**Preconditions:**  
**Constraints:** W(S, Self:Action)

**Speech-Act(S:Agent, H:Agent)**  
**Preconditions:** Attend(H, S), Attend(S, H)

As we look deeper into the hierarchy, we find that the leaf speech acts embody distinctions which may be closely tied to the language. Any language that



|  
Informref(S:Agent, H:Agent, P(x), x)

In this fragment of the speech act hierarchy, the three Informing acts are siblings, as is Speculation. This is not a complete subtree, since there may be other Representative-Acts. We suggest here that Speculation communicates that some proposition is possible, in the sense that it is not known but is consistent with what is known. Such an act should have an effect on attention, but this is simply an artifact of the communication process rather than something formalized in the act's definition. An Inform act requires that the speaker believe the proposition and intends that the hearer believe the same proposition. Its body is the reflexive Gricean intention, so that any way of achieving recognition of this intention will satisfy the Inform. Practically speaking several linguistic rules generate Informs. Informif is weaker; it doesn't represent whether the communicated proposition is P or  $\neg P$ . Its body can be satisfied by Inform(S, H, P) and by Inform(S, H,  $\neg P$ ). An Informref can be satisfied by Inform(S, H, P) where the appropriate variable is bound.

**Speculate(S:Agent, H:Agent, P)**  
**Preconditions:**  $\neg$  Knowif(S, P)  
**Constraints:** B(S, Possible(P))  
**Body:** B(H, W(S, B(H, Possible(P))))  
**Effects:** B(H, Possible(P))

**Inform(S:Agent, H:Agent, P)**  
**Constraints:** B(S, P)  
**Body:** B(H, W(S, B(H, P)))  
**Effects:** B(H, P)

**Informif(S:Agent, H:Agent, P)**  
**Preconditions:** Knowif(S, P)  
**Constraints:**



Body: B(H, W(S, Knowif(H, P)))  
 Effects: Knowif(H, P)

Informref(S:Agent, H:Agent, P(x))  
 Preconditions: Knowref(S, P(x))  
 Constraints: B(S, P)  
 Body: B(H, W(S, B(H, P)))  
 Effects: B(H, P)

Questioning acts are as closely related to representatives as to directives, as we found in Chapter Three. We distinguished three kinds of questions: yes/no questions, indicating ignorance of the truth value of a proposition, disjunctive questions, which specify a set of alternative values for a variable, and wh questions, which mark a variable but leave the set of values unspecified.

```

Ask-Act(S, H, P)-----|
|                         |
|                         | Askif(S, H, P)
|                         |
|                         |
| Askor(S, H, P, ?Values)
|
|
Askwh(S, H, P, ?Variable)
  
```

People plan Ask-Acts in order to cause Representative-Acts by other agents. This only works under the constraint that the other agent is actually able to perform the act, and it occurs any time that agent is convinced by the first one that the act is wanted. This Ask-Act is a rough abstraction of the three more specific types of questions discussed in Chapter Three. Askif corresponds to yes/no questions, Askor to disjunctive questions, and Askwh to wh-questions. Each of these acts requires explicitly that the hearer have the belief that would answer the question,

and that the speaker want to have it. Rhetorical questions do not have this requirement; they could be incorporated into the hierarchy as cousins of these acts, or they can be derived by extended reasoning on each occasion. Didactic questions, in which the questioner knows the answer and the hearer may not, are also cousins of these acts.

**Ask-Act(S:Agent, H:Agent, P)**

Preconditions:

Constraints: Able(H, Representative-Act(H, P))

Body: DO(S, B(H, W(S, Representative-Act(H, S, P))))

Effects: Representative-Act(S, H, P)

**Askif(S:Agent, H:Agent, P)**

Preconditions:

Constraints: Able(H, Informif(H, S, P)),

Knowif(H, P),

W(S, Knowif(S, P))

Body: DO(S, B(H, W(S, Informif(H, S, P))))

Effects: Informif(H, S, P)

**Askwh(S:Agent, H:Agent, P(x))**

Preconditions:

Constraints: Able(H, Informref(H, S, P(x))),

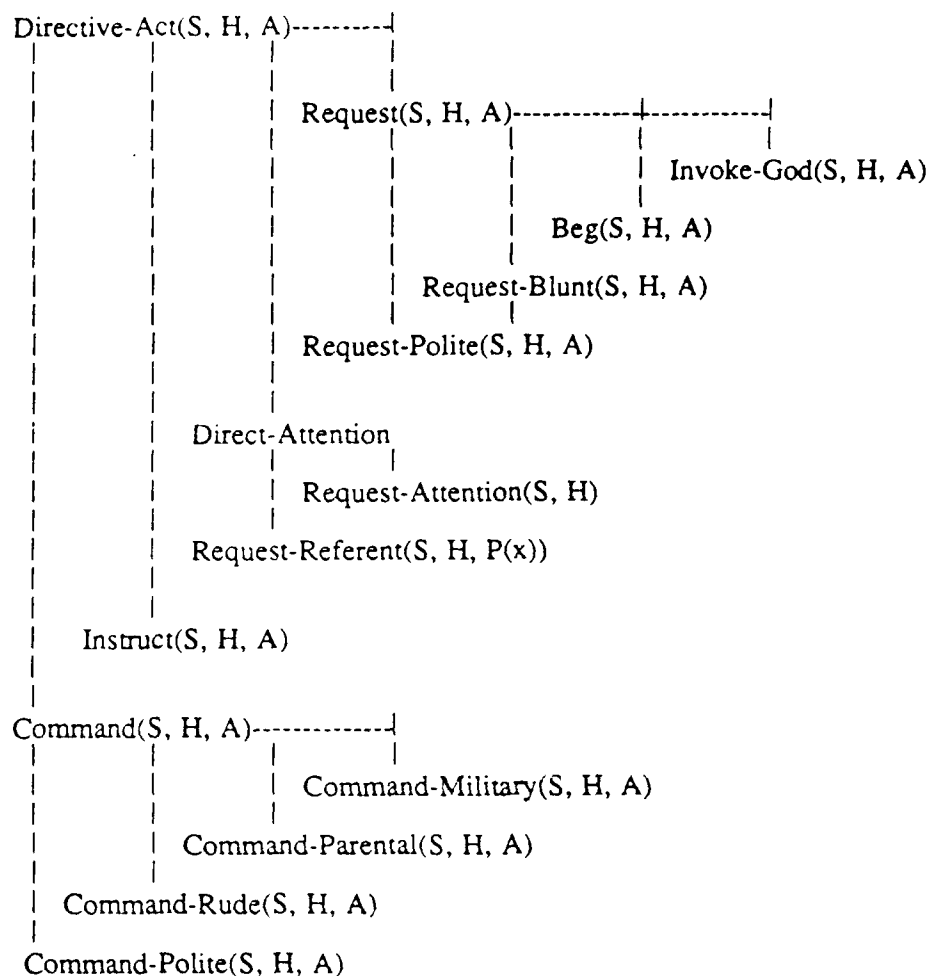
Knowref(H, P(x)),

W(S, Knowref(S, P(x)))

Body: DO(S, B(H, W(S, Informref(H, S, P(x)))))

Effects: Informref(H, S, P(x))

Directive acts come in surprising variety. Commands are based on a power relationship, but may be polite or a part of a specialized sublanguage like military commands. Instructions too constitute a sublanguage, but assume that the goal of the task is one that the hearer already has for some reason. Directing someone's attention is qualitatively much different, but probably more fundamental. Requests are the most familiar, but even here we must add the more colorful acts of begging and of invoking a deity.



The most abstract Directive-Act has only the intended effect that the hearer do some action. Eventually a distinction will have to be made between hearers of a directive act and actual addressees who are intended to carry out the directive. Commands all require an authority relation. We will assume that this entails a belief that something bad will happen to the hearer if the command is not heeded,

but we will not attempt to formalize this condition. The polite command is used by a speaker who wishes not to offend the hearer, but who nonetheless intends the action to be done. Request-Attention is equivalent to a Request in which the Requested action is to pay attention. Request-Referent is really a referring action: the speaker wants the hearer to identify internally the described referent [Perrault 78].

**Directive-Act(S, H, A)**

Body: B(H, W(S, Do(H, A)))

Effects: DO(H, A)

**Command(S, H, A)**

Preconditions: SUPERIOR(S, H)

**Command-Polite(S, H, A)**

Constraints: W(S, ~ Offend(S, H))

**Command-Rude(S, H, A)**

**Command-Parental(S, H, A)**

Preconditions: Parent(S, H)

**Command-Military(S, H, A)**

Preconditions: Superior-Officer(S, H)

**Request-Attention(S, H, Attend(H, S))**

Effect: Attend(H, S)

**Request-Referent(S, H, P(x))**

Knowref(H, P(x))

**Request(S, H, A)**

Constraints: Able(H, DO (H, A)), W(S, Effects(A))

**Request-Polite(S, H, A)**

Preconditions: W(S, ~ Offend(S, H))

**Request-Rude(S, H, A)**

**Beg(S, H, A)**



Accept(S:Agent, H:Agent, A)  
 Preconditions: Offer(H, S, A)  
 Constraints: W(S, A)  
 Effects: Do(H, A)

Here is a pair of Declarative-Acts, generally realized by explicit performative utterances:

Resign(S, H, Position)  
 Preconditions:  
 Constraints: Holds(S, Position), Employer(H, S)  
 Effects:  $\neg$  Holds(S, Position)

Fire(S, H, Position)  
 Preconditions:  
 Constraints: Holds(H, Position), Employer(S, H)  
 Effects:  $\neg$  Holds(H, Position)

### 5.3. Some Simple Examples

Each implicature schema provides implicatures that can serve to filter out impossible interpretations. We now consider examples based on the different categories of implicatures.

#### 5.3.1. Speech Act Preconditions

Suppose that Pat and Sandy share an office. When Sandy returns from a meeting, Pat says,

(71) Your husband called.

We model a possible interpretation of the utterance as an Inform of the literal content. This interpretation is shown below:

**Inform(Pat, Sandy, Phone(Sandy, Husband(Sandy)))**  
**Preconditions: Knowif(Pat, Phone(Sandy, Husband(Sandy))),**  
**Attend(Sandy, Pat), Attend(Pat, Sandy)**  
**Constraints: Able(Pat, Self:Action), W(Pat, Self:Action),**  
**B(Pat, Phone(Sandy, Husband(Sandy)))**  
**Body: B(Sandy, W(Pat, Knowif(Sandy, Phone(Sandy, Husband(Sandy))))**  
**Effects: K(Sandy, Phone(Sandy, Husband(Sandy)))**

This is simply the schema for an Inform act, including its inherited conditions, with the two agents and the proposition substituted. Now, suppose Sandy's beliefs include

Attend(Sandy, Pat)  
 Attend(Pat, Sandy)

and that Sandy believes Pat shares these beliefs. For Sandy to accept the Inform interpretation, several implicatures must be consistent. First, the hearer must believe that the speaker believes that the preconditions hold. This includes the speaker's knowing the fact, and the observation conditions. In this case the observation conditions are known and the Knowif is implicated.

preconditions hold  
 B(Sandy, B(Pat, Knowif(Pat, Phone(Sandy, Husband(Sandy))))  
 B(Sandy, B(Pat, Attend(Sandy, Pat)))  
 B(Sandy, B(Pat, Attend(Pat, Sandy)))

The hearer must believe that the speaker may believe the constraints hold. Both the speaker's wanting the action and the speaker's believing the proposition are implicated.

constraints hold  
 B(Sandy, B(Pat, W(Pat, Self:Action)))  
 B(Sandy, B(Pat, B(Pat, Phone(Sandy, Husband(Sandy))))

The hearer believes that the speaker intended the effects, namely, that the hearer believe the fact. This is implicated.

effects intended  
 B(Sandy, W(Pat, B(Sandy, Phone(Sandy, Husband(Sandy))))))

The hearer believes the speaker believes that the effects were not already true, namely, that Sandy didn't know her husband called. This is implicated.

effects didn't hold  
 B(Sandy, B(Pat, ~B(Sandy, Phone(Sandy, Husband(Sandy))))))

Sandy already believes that the observation conditions hold, but has no beliefs in which the phone call appears. The beliefs about the Inform act itself and about the phone call are therefore implicated, if the Inform interpretation is accepted. Since no contradictions arise, the Inform interpretation is acceptable.

We contrast the outcome in this context with that in a related context with some important differences. Suppose that when Sandy arrived at the office, Pat had his back to her and was facing Liz. Sandy's context would then be  
 Attend(Pat, Liz)

~ Attend(Sandy, Pat)  
 ~ Attend(Pat, Sandy)

The beliefs about the phone call would still be possible implicatures. However, since the observation conditions fail, the Inform(Pat, Sandy..) interpretation is not acceptable. Sandy might recognize an Inform(Pat, Liz..) act instead, with its own



implicatures. The precondition schema allows the Inform(Pat, Sandy...) interpretation to be eliminated by identifying parts of the context which are inconsistent with the interpretation.

### 5.3.2. Negations of Effects

In this example, an interpretation has an effect which is already true and is implausible for that reason.

Dana is going out the door.

(72) Dana: I have to go and pick up the kids.  
Sandy: Oh, so the class is done at 6.

If Sandy were Informing Dana that class ends at 6, the interpretation would be

```
Inform(Sandy, Dana, Precedes(Now, End_Time(Class007)))
  Preconditions: Attend(Sandy, Pat), Attend(Pat, Sandy)
  Constraints: W(Sandy, Self:Action),
             B(Sandy, Equals(6, End_Time(Class007)))
  Body: B(Dana, W(Sandy, K6if(Dana,
                             Equals(6, End_Time(Class007)))))
  Effects: B(Dana, Equals(6, End_Time(Class007)))
```

Dana believes that Sandy believes

```
K6if(Dana, Equals(6, End_Time(Class007)))
Attend(Sandy, Dana), Attend(Dana, Sandy)
B(Dana, Equals(6, End_Time(Class007)))
```

However, the implicatures of the Inform act would be

```
preconditions hold
B(Dana, B(Sandy, Attend(Sandy, Dana)))
B(Dana, B(Sandy, Attend(Dana, Sandy)))
```

constraints hold  
 B(Dana, B(Sandy, W(Sandy, Self:Action)))  
 B(Dana, B(Sandy, B(Sandy, Equals(6, End\_Time(Class007)))))

effects intended  
 B(Dana, W(Sandy, B(Dana, Equals(6, End\_Time(Class007)))))

effects do not hold  
 B(Dana, B(Sandy, ~B(Dana, Equals(6, End\_Time(Class007)))))

Dana, the putative recipient, already believes the information. The speech act's effect is therefore true, and the implicature that this effect doesn't already hold is false. The Inform interpretation is eliminated. Sandy is really asking for confirmation of a fact inferred from the first utterance, as the word "so" indicates.

### 5.3.3. Intended Effects

In the next example, the effect schema yields the contradiction. Sandy is relating an old sticky situation to a new boss, Jan.

(73) Jan: What did you do when he insisted?  
 Sandy: I quit.

In some contexts, saying "I quit" is to resign from one's job. Here a Resign interpretation would look like this:

**Resign(Sandy, Jan, Position27)**  
**Preconditions:** Attend(Jan, Sandy), Attend(Sandy, Jan)  
**Constraints:** Holds(Sandy, Position27), Employer(Jan, Sandy)  
**Effects:** ~Holds(Sandy, Position27)

We model Jan as believing Sandy shares the following beliefs:

Attend(Jan, Sandy), Attend(Sandy, Jan)  
 Holds(Sandy, Position27), Employer(Jan, Sandy)

W(Sandy, Holds(Sandy, Position27))

The speech act has these implicatures:

preconditions hold

B(Jan, B(Sandy, Channel(Sandy, Jan)))

B(Jan, B(Sandy, Attend(Jan, Sandy)))

B(Jan, B(Sandy, Attend(Sandy, Jan)))

constraints hold

B(Jan, B(Sandy, Holds(Sandy, Position27)))

B(Jan, B(Sandy, Employer(Jan, Sandy)))

effects intended

B(Jan, W(Sandy, ~ Holds(Sandy, Position27)))

effects do not hold

B(Jan, B(Sandy, Holds(Sandy, Position27)))

The Resign interpretation arises under a present-tense reading of the utterance, and could be rejected by a temporal module favoring continuity of tense. However, the interpretation can be eliminated by its implicatures too. Sandy is known to want to keep the job. In our logic it is not possible to want both a state and its negation. Thus the effect of the speech act is unintended, and the implicature that it is intended yields a contradiction. Jan will not believe that Sandy is actually resigning.

#### 5.3.4. Constraints

Here is a similar example in which a constraint fails. Sandy storms home from work, and announces to Dana

(74) I quit!

The Resign act is

**Resign(Sandy, Dana, Position27)**

**Preconditions:** Attend(Dana, Sandy), Attend(Sandy, Dana)

**Constraints:** Holds(Sandy, Position27), Employer(Dana, Sandy)

**Effects:** ~ Holds(Sandy, Position27)

Dana believes that Sandy shares these beliefs:

Attend(Dana, Sandy), Attend(Sandy, Dana)

Holds(Sandy, Position27), Employer(Jan, Sandy), ~ Employer(Dana, Sandy)

Let us suppose that Dana is agnostic about Sandy's desire for the job. The implicatures are

preconditions hold

B(Dana, B(Sandy, Attend(Dana, Sandy)))

B(Dana, B(Sandy, Attend(Sandy, Dana)))

constraints hold

B(Dana, B(Sandy, Holds(Sandy, Position27)))

B(Dana, B(Sandy, Employer(Dana, Sandy)))

effects intended

B(Dana, W(Sandy, ~ Holds(Sandy, Position27)))

effects do not hold

B(Dana, B(Sandy, Holds(Sandy, Position27)))

The second constraint-based implicature contradicts Dana's belief that Sandy knows Dana isn't boss, and therefore makes the Resign interpretation unacceptable. Dana may conclude that Sandy has already quit, or that Sandy intends to quit.

Each of our implicature schemas has proven itself useful in eliminating speech act interpretations that are inconsistent with context. When the interpretations are not eliminated, the implicatures serve as new and useful information to the hearer. We next consider the overall interpretation process.

## 6. Two Constraints Integrated

Chapters Two and Three showed how to compute a set of possible speech act interpretations incrementally, from conventions of language use. Chapters Four and Five showed how plan reasoning, which motivates the conventions, can be used to develop further interpretations and to eliminate implausible interpretations. In Chapter Six we show how the components can be integrated with each other to handle the full range of speech acts.

The interface between the linguistic and implicature components is very simple. The linguistic component yields a set of speech act interpretations, and the implicature computation takes this set as input and acts as a filter on it. The implicature computation therefore yields a reduced set of speech act interpretations. *This reduced set of interpretations can then be input extended plan reasoning or accepted, under criteria to be specified below.* The fact that speech acts are explicitly represented the association of detailed linguistic patterns with detailed patterns of propositions, and the interfaces among the components are simply sets of these explicit representations.

The overall process is that along with the usual incremental linguistic processes, we build up and merge hypotheses about speech act interpretations. The resulting interpretations are passed to the implicature module. The conversational implicatures are computed, discounting interpretations if they are in conflict with contextual knowledge. The interpretations remaining may be passed to extended

reasoning. If a plausible, non-contradictory interpretation remains, it can be accepted. Allen-style plan reasoning is invoked to identify the speech act only if remaining ambiguity interferes with planning or if no completely plausible interpretations remain.

We do not address the control issues raised by extended reasoning in any comprehensive way. Our method of speech act interpretation avoids extended reasoning where possible. It requires only that interpretations proposed by extended reasoning fit the linguistic module's constraints, and that the implicatures of any final interpretation be consistent. It does not require that there be a final interpretation.

### **6.1. *Interaction of the Constraints***

The linguistic computation constrains plan reasoning by providing the input. The final interpretation must fall within the range of the input. In more concrete terms, it is as if the observed act were asserted to be equal to some subset of the disjunction output by the linguistic module. Further processing must be consistent with this equality.

Recall that the linguistic rules control ambiguity: because the right hand side of the rule must express *all* the possibilities for this pattern, a single rule can limit the range of interpretations sharply. Consider

- (75) a: I hereby inform you that it's cold in here.  
       b: It's cold in here.

The explicit performative rules, triggered by "hereby" and by a performative verb in the appropriate syntactic context, allow for only an explicit performative interpretation of sentence (a). The linguistic module yields only the Inform interpretation, and subsequent processing may give further detail to the Inform or render it implausible. However, it cannot propose a non-Inform interpretation, since this would fall outside the range indicated by the linguistic module. By contrast, the declarative rule proposes two speech acts for (b), the Inform and the abstract SpeechAct. Since the SpeechAct encompasses many subtypes, it allows the plan reasoner to identify other interpretations for (b).

The plan reasoning phase constrains the results of the linguistic computation by eliminating interpretations, and reinterpreting others. In a context where the speaker and hearer mutually believe that it's cold, the Inform interpretation is filtered out by implicature checking. For sentence (a) this would leave no plausible interpretations (so that the system must ask what was meant or reason about possible misconceptions.) For sentence (b) plan reasoning would eliminate the linguistic module's Inform interpretation. In this context it would perform the extended reasoning we discussed earlier, in which the utterance is identified as a Request. The implicature check eliminates some interpretations, and the extended reasoning refines the more abstract interpretation by identifying a likely interpretation that specializes the abstract one.

## 6.2. An Extended Example

As an illustration of the combined action of the linguistic and plan reasoning components, we consider how two related sentences are interpreted, in each of two related contexts. The sentences are "Can you speak Spanish?", and "Can you speak Spanish, please?". They differ only in the addition of the word "please". The contexts differ in Suzanne's model as hearer, of whether Mrs. de Prado believes Suzanne can speak Spanish. Since implicature checking will filter interpretations differentially according to context, the results are sensitive to both linguistic and contextual variation.

In the first context, Suzanne is at the Spanish consulate, doing her paperwork for a Fulbright scholarship year in Spain. Mrs. de Prado, the representative, asks, "Can you speak Spanish?" Suppose that Suzanne has previously declared her fluency in Castilian. Her belief space is

### Context One.

```
MB(Suzanne, Mrs. de Prado, Attend(P, S))
MB(Suzanne, Mrs. de Prado, Attend(S, P))
...
MB(Suzanne, Mrs. de Prado, Able(Suzanne,
    Use-Language(Suzanne, Spanish)))
MB(Suzanne, Mrs. de Prado, Able(Suzanne, Informif(Suzanne,
    Mrs. de Prado, Able(Suzanne, Use(Suzanne, Spanish))))
MB(Suzanne, Mrs. de Prado, B(Mrs. de Prado, Able(Suzanne,
    Use-Language(Suzanne, Spanish))))
```

The second context is similar except that Suzanne has not previously declared her fluency in Castilian. She may hold these beliefs:



**Context Two.**

MB(Suzanne, Mrs. de Prado, Attend(P, S))

MB(Suzanne, Mrs. de Prado, Attend(S, P))

Able(Suzanne, Use-Language(Suzanne, Spanish))

MB(Suzanne, Mrs. de Prado, Able(Suzanne, Informif(Suzanne,  
Mrs. de Prado, Able(Suzanne, Use(Suzanne, Spanish))))

MB(Suzanne, Mrs. de Prado, ~Knowif(Mrs. de Prado,  
Able(Suzanne, Use-Language(Suzanne, Spanish))))

In both cases she knows that she can speak Spanish, but only in the first case does she believe that Mrs. de Prado knows this. We will consider how the utterance "Can you speak Spanish?" fares in each of these contexts. Having seen the effects of context on its interpretation, we can then compare "Can you speak Spanish, please?" and how context affects it.

**6.2.1. Can you speak Spanish? -- Context One**

For this utterance, we recall that the linguistic computation yields three interpretations:

((REQUEST-ACT AGENT Mrs. de Prado  
HEARER Suzanne  
ACTION (USE-LANGUAGE AGENT Suzanne  
LANG ls1)))

(ASKIF AGENT Mrs. de Prado  
HEARER Suzanne  
PROP (ABLE-STATE AGENT Suzanne  
ACTION (USE AGENT Suzanne  
OBJECT ls1)))

(SPEECH-ACT AGENT Mrs. de Prado)  
HEARER Suzanne))

How do these interpretations fare under implicature calculation? The Request interpretation's full description, including inherited conditions, is

Request(Mrs. de Prado, Suzanne, Use(Suzanne, Spanish))  
 Preconditions: Attend(Suzanne, Mrs. de Prado),  
                 Attend(Mrs. de Prado, Suzanne)  
 Constraints: Able(Mrs. de Prado, Self:Action),  
                 W(Mrs. de Prado, Self:Action),  
                 W(Mrs. de Prado, Effects(Self)),  
                 Able(Suzanne, Use-Language(Suzanne, Spanish))  
 Effects: DO(Suzanne, Use(Suzanne, Spanish))

Our algorithm checks the Request's implicatures in Context One. The complete list is:

preconditions hold:

B(Suzanne, B(Mrs. de Prado, Attend(Suzanne, Mrs. de Prado)))  
 B(Suzanne, B(Mrs. de Prado, Attend(Mrs. de Prado, Suzanne)))

constraints hold:

B(Suzanne, B(Mrs. de Prado, Able(Mrs. de Prado, Request(...))))  
 B(Suzanne, B(Mrs. de Prado, W(Mrs. de Prado, Request(...))))  
 B(Suzanne, B(Mrs. de Prado, W(Mrs. de Prado, Effects(Use(Suzanne, Spanish)))))  
 B(Suzanne, B(Mrs. de Prado, Able(Suzanne, Use-Language(Suzanne, Spanish))))

effects intended:

B(Suzanne, W(Mrs. de Prado, DO(Suzanne, Use(Suzanne, Spanish))))

effects do not hold:

B(Suzanne, B(Mrs. de Prado, ~DO(Suzanne, Use(Suzanne, Spanish))))

All of the implicatures above are consistent with what Suzanne knows, so the interpretation is not filtered out. The preconditions are explicit in the context, and Mrs. de Prado's general ability to perform speech acts is background information. Suzanne's ability and not already speaking Spanish follow from context and are therefore consistent but not implicated. Mrs. de Prado's wanting to Request, wanting Suzanne to speak Spanish, and wanting the effects of this action are new information and are the most important aspect of the request. They are conclusions which will be asserted when this interpretation is ultimately accepted.

This interpretation thus passes from the linguistic module through implicature checking and yields significant new information.

The question interpretation's full description is:

**Askif(Mrs. de Prado, Suzanne, Able(Suzanne, Use(Suzanne, Spanish)))**  
**Preconditions: Attend(Suzanne, Mrs. de Prado),**  
**Attend(Mrs. de Prado, Suzanne)**  
**Constraints: Able(Mrs. de Prado, Self:Action), W(Mrs. de Prado, Self:Action)**  
**~ Knowif(Mrs. de Prado,**  
**Able(Suzanne, Use(Suzanne, Spanish)))**  
**Able(Suzanne, Informif(Suzanne, Mrs. de Prado,**  
**Able(Suzanne, Use(Suzanne, Spanish))))**  
**Knowif(Suzanne, Able(Suzanne,**  
**Use(Suzanne, Spanish)))**  
**W(Mrs. de Prado, Knowif(Mrs. de Prado, Able(Suzanne,**  
**Use(Suzanne, Spanish))))**  
**Effects: Informif(Suzanne, Mrs. de Prado, Able(Suzanne, Use(Suzanne, Spanish)))**

Its implicatures are

```
preconditions hold:
```

B(Suzanne, B(Mrs. de Prado, Attend(Suzanne, Mrs. de Prado)))  
B(Suzanne, B(Mrs. de Prado, Attend(Mrs. de Prado, Suzanne)))

constraints hold:

B(Suzanne, B(Mrs. de Prado, Able(Mrs. de Prado, Askif(...)))  
B(Suzanne, B(Mrs. de Prado, W(Mrs. de Prado, Askif(...)))  
B(Suzanne, B(Mrs. de Prado, ~ Knowif(Mrs. de Prado,  
Able(Suzanne, Use(Suzanne, Spanish))))

B(Suzanne, B(Mrs. de Prado, Able(Suzanne, Informif(Suzanne, Mrs. de Prado, Able(Suzanne, Use(Suzanne, Spanish))))))

B(Suzanne, B(Mrs. de Prado, Knowif(Suzanne,  
Able(Suzanne, Use(Suzanne, Spanish))))

B(Suzanne, B(Mrs. de Prado, W(Mrs. de Prado, Knowif(Mrs. de Prado, Able(Suzanne, Use(Suzanne, Spanish))))))

**effects intended:**

B(Suzanne, W(Mrs. de Prado, Informif(Suzanne, Mrs. de Prado,  
Able(Suzanne, Use(Suzanne, Spanish))))))

effects do not hold:

B(Suzanne, B(Mrs. de Prado, ~ Informif(Suzanne, Mrs. de Prado, Able(Suzanne, Use(Suzanne, Spanish))))

In Context One, the third constraint fails. Since Suzanne believes Mrs. de Prado knows Suzanne can speak Spanish, it is inconsistent for Suzanne to believe that Mrs. de Prado believes herself ignorant of the fact. Therefore the Ask interpretation is eliminated in this context. The third interpretation is the abstract one:

**Speech-Act(Mrs. de Prado:Agent, Suzanne:Agent)**  
**Preconditions:** Attend(Suzanne, Mrs. de Prado),  
 Attend(Mrs. de Prado, Suzanne)  
**Constraints:** Able(Mrs. de Prado, Self:Action),  
 W(Mrs. de Prado, Self:Action)

Its implicatures are:

preconditions hold:

B(Suzanne, B(Mrs. de Prado, Attend(Suzanne, Mrs. de Prado)))  
 B(Suzanne, B(Mrs. de Prado, Attend(Mrs. de Prado, Suzanne)))

constraints hold:

B(Suzanne, B(Mrs. de Prado, Able(Mrs. de Prado, Speech-Act(...))))  
 B(Suzanne, B(Mrs. de Prado, W(Mrs. de Prado, Speech-Act(...))))

These are all consistent, although Mrs. de Prado's intent to perform a speech act is new information and would be implicated. They are consistent with the Request interpretation. In Context One we have for "Can you speak Spanish?" two consistent interpretations, the Request and the Speech-Act. The two are also consistent with each other.

### 6.2.2. Can you speak Spanish? -- Context Two

Let us now take the same set of interpretations, and consider what happens to them in the second context. In this context Suzanne believes that Mrs. de Prado is ignorant of her Spanish skills. The effect of this difference is to eliminate the

Request interpretation rather than the Ask.

The Request interpretation's implicatures are, again:

preconditions hold:

B(Suzanne, B(Mrs. de Prado, Attend(Suzanne, Mrs. de Prado)))

B(Suzanne, B(Mrs. de Prado, Attend(Mrs. de Prado, Suzanne)))

constraints hold:

B(Suzanne, B(Mrs. de Prado, Able(Mrs. de Prado, Request(...))))

B(Suzanne, B(Mrs. de Prado, W(Mrs. de Prado, Request(...))))

B(Suzanne, B(Mrs. de Prado, W(Mrs. de Prado, Effects(Use(Suzanne, Spanish)))))

B(Suzanne, B(Mrs. de Prado, Able(Suzanne, Use-Language(Suzanne, Spanish))))

effects intended:

B(Suzanne, W(Mrs. de Prado, DO(Suzanne, Use(Suzanne, Spanish))))

effects do not hold:

B(Suzanne, B(Mrs. de Prado, ~DO(Suzanne, Use(Suzanne, Spanish))))

A constraint on the Request interpretation fails. It requires that the speaker believe the hearer can perform the action being requested. In this context Suzanne does not believe Mrs. de Prado has that belief. Thus the Request interpretation is filtered out by the implicature check in this context.

The Ask act's constraints now hold. Mrs. de Prado can sincerely ask because she doesn't know the answer to her question. The Ask interpretation is not eliminated as it was before.

preconditions hold:

B(Suzanne, B(Mrs. de Prado, Attend(Suzanne, Mrs. de Prado)))

B(Suzanne, B(Mrs. de Prado, Attend(Mrs. de Prado, Suzanne)))

constraints hold:

B(Suzanne, B(Mrs. de Prado, Able(Mrs. de Prado, Askif(...))))

B(Suzanne, B(Mrs. de Prado, W(Mrs. de Prado, Askif(...))))

B(Suzanne, B(Mrs. de Prado, ~Knowif(Mrs. de Prado,  
Able(Suzanne, Use(Suzanne, Spanish)))))

B(Suzanne, B(Mrs. de Prado, Able(Suzanne, Informif(Suzanne, Mrs. de Prado,

Able(Suzanne, Use(Suzanne, Spanish))))))  
 B(Suzanne, B(Mrs. de Prado, Knowif(Suzanne,  
 Able(Suzanne, Use(Suzanne, Spanish))))))  
 B(Suzanne, B(Mrs. de Prado, W(Mrs. de Prado, Knowif(Mrs. de Prado,  
 Able(Suzanne, Use(Suzanne, Spanish))))))

effects intended:

B(Suzanne, W(Mrs. de Prado, Informif(Suzanne, Mrs. de Prado,  
 Able(Suzanne, Use(Suzanne, Spanish))))))

effects do not hold:

B(Suzanne, B(Mrs. de Prado, ~Informif(Suzanne, Mrs. de Prado,  
 Able(Suzanne, Use(Suzanne, Spanish))))))

The new implicatures are the second and sixth based on constraints, and the intended effect. They convey the speaker's intent to ask a question.

The Speech-Act's implicatures all go through as before. Thus just the Ask and Speech-Act remain.

It is left as an exercise to the reader to show that if Mrs. de Prado believed Suzanne could not speak Spanish, both Request and Ask would be eliminated.

### 6.2.3. Can you speak Spanish, please? -- Context One

We now return to Context One, in which Suzanne believes that Mrs. de Prado knows she speaks Spanish. The sentence "Can you speak Spanish, please?" has these possible interpretations:

((REQUEST-ACT AGENT Mrs. de Prado  
 HEARER Suzanne  
 ACTION (USE-LANGUAGE AGENT Suzanne  
 LANG ls1)))

((DIRECTIVE-ACT AGENT Mrs. de Prado  
 HEARER Suzanne)

The Request interpretation is just the same as for the first sentence. We have

already seen that the Request interpretation is consistent in Context One. The

Directive act is similar:

**Directive-Act**(Mrs. de Prado, Suzanne, Use(Suzanne, Spanish))  
**Preconditions:** Attend(Suzanne, Mrs. de Prado),  
Attend(Mrs. de Prado, Suzanne)  
**Constraints:** Able(Mrs. de Prado, Self:Action),  
W(Mrs. de Prado, Self:Action)  
**Effects:** DO(Suzanne, Use(Suzanne, Spanish))

The implicatures are a subset of those for the Request.

preconditions hold:

B(Suzanne, B(Mrs. de Prado, Attend(Suzanne, Mrs. de Prado)))  
B(Suzanne, B(Mrs. de Prado, Attend(Mrs. de Prado, Suzanne)))

constraints hold:

B(Suzanne, B(Mrs. de Prado, Able(Mrs. de Prado, Directive-Act(...))))  
B(Suzanne, B(Mrs. de Prado, W(Mrs. de Prado, Directive-Act(...))))

effects intended:

B(Suzanne, W(Mrs. de Prado, DO(Suzanne, Use(Suzanne, Spanish))))

effects do not hold:

B(Suzanne, B(Mrs. de Prado, ~DO(Suzanne, Use(Suzanne, Spanish))))

The Directive act is therefore also consistent. The implicatures which provide new information are those involving Mrs. de Prado's intentions to perform a Directive act and to have Suzanne speak Spanish. These are abstractions of those for the Request, and in the presence of an authority relation, the act could be specialized to a polite command. In Context One both the Request and the Directive-Act are consistent, and they are consistent with each other.

#### 6.2.4. Can you speak Spanish, please? -- Context Two

In Context 2, the Request act fails as it did previously in this context, because Mrs. de Prado does not believe that Suzanne can speak Spanish. The Directive-Act, which does not require that the agent is able to do the action, goes through.

preconditions hold:

B(Suzanne, B(Mrs. de Prado, Attend(Suzanne, Mrs. de Prado)))

B(Suzanne, B(Mrs. de Prado, Attend(Mrs. de Prado, Suzanne)))

constraints hold:

B(Suzanne, B(Mrs. de Prado, Able(Mrs. de Prado, Directive-Act(...))))

B(Suzanne, B(Mrs. de Prado, W(Mrs. de Prado, Directive-Act(...))))

effects intended:

B(Suzanne, W(Mrs. de Prado, DO(Suzanne, Use(Suzanne, Spanish))))

effects do not hold:

B(Suzanne, B(Mrs. de Prado, ~DO(Suzanne, Use(Suzanne, Spanish))))

The precondition-based propositions hold in this context, as does the one based on negations of effects. The intentions are implicated. Thus we have the Directive interpretation only.

#### 6.2.5. Comparison

The comparison between the two sentences in the various contexts is summarized in the figure below. A's indicate acceptable interpretations, and X's contradictions.

In the first context, where Suzanne is known to speak Spanish, the question interpretation is eliminated for the first utterance. This leaves the Request and the Speech-Act. Since the Request specializes the Speech-Act, there is no contradiction to be resolved. The more abstract act allows for the possibility that



Can you speak Spanish?    B(H, B(S, Able...))    B(H, ~Knowif(S, Able...))    B(H, B(S, ~Able...))

Askif	X	A	X
Request	A	X	X
Speech-Act	A	A	A

Can you speak Spanish, please?

Request	A	X	X
Directive-Act	A	A	A

another of its specializations has occurred, but restricts any extended plan reasoning to interpretations in this range. For the second utterance, the Request also specializes the Directive-Act.

In the second context, when Suzanne believes Mrs. de Prado to be ignorant of her Spanish ability, the first utterance cannot be a Request. There remains the Askif specializing the Speech-Act. The second utterance likewise cannot be a Request, leaving the Directive-Act.

The third context was mentioned in passing above. In the third context, where Suzanne is known not to speak Spanish, the first utterance retains only the Speech-Act. The second utterance retains only the Directive-Act. Thus implicatures screen out different interpretations according to the context, by pinpointing the relevant portions of that context.

The extended inference process can never yield an Ask interpretation for the second utterance, because the input from the linguistic module is already too

narrow. Any interpretation of the utterance must fall within the range of interpretations output from the linguistic module, and while the first utterance yields an abstract speech act with many possible specializations, the second utterance does not. The Ask interpretation can therefore never be proposed by subsequent processing, unless an error is postulated which accounts for the discrepancy. Overall we thus have an interpretation process which is sensitive both to linguistic variation and contextual variation.

### **6.3. Extended Reasoning**

Extended reasoning is still necessary in some cases. The agent may face an unfamiliar speech act, and work it out based on context or the set of rejected interpretations. The agent may need to reduce remaining ambiguity, for planning purposes. There may be a misconception on the part of the speaker or the hearer, including the hearer's model of the speaker. In the long term, extended plan reasoning also makes it possible for some new conventions of language use to develop. It therefore becomes necessary to provide an interface between the plan reasoning we have been discussing and extended reasoning.

#### **6.3.1. The Interface**

The interface between the short and extended reasoning is a simple one. The set of interpretations resulting from the implicature computation, as described above, can be accepted as the interpretation if all interpretations in the set are mutually consistent, and if they need not be distinguished for planning purposes. These

conditions are trivially true when there is only one interpretation, or when the interpretations abstract each other.

The alternative is to invoke extended plan reasoning. Extended reasoning may be invoked for ambiguity resolution if multiple interpretations remain and these interpretations need to be distinguished for planning purposes. Extended reasoning may be invoked to derive further speech act interpretations if the set contains no interpretations, or only very abstract actions such as Speech-Acts or Representative-Acts. In either case it is assumed that the act's type restrictions have already been asserted. We summarize this information below.

If there are no remaining interpretations, or abstract ones only:  
 Invoke plan reasoning to derive a new interpretation

If there is one interpretation, or one and abstractions of it:  
 Accept the most specific interpretation and its implicatures

If there are several remaining interpretations, distinguished by the planner,  
 Invoke plan reasoning to disambiguate

Let's consider a few examples.

When no interpretations remain, plan reasoning can infer novel interpretations. In a foreign restaurant, the waiter says to you "your soup, please," with bowl in hand. The linguistic module restricts the interpretation to a request, but there is no appropriate action to be requested. There are no appropriate interpretations. Plan reasoning must be invoked on the available information to determine that the waiter is offering you the soup. One explanation for the utterance is that he is

using "please" as an honorific.

Another example of this sort is "Can you speak Spanish, please?", from a person who believes the hearer cannot speak Spanish. It cannot be a sincere directive, but reasoning about plans may yield an attempt to humiliate, or some spy-story plot.

Examples of a single interpretation, and one with an additional abstraction, can be found among the Spanish utterances. The second sentence and second context yielded only the Directive interpretation. This may be accepted, and its implicatures asserted. A single, unambiguous interpretation arises rarely, since there is usually a more abstract interpretation present as well. This situation arises in the Spanish example's first context, for both sentences. For "Can you speak Spanish?" the possibilities are the abstract speech act and the Request; the hearer can accept the Request and its implicatures. For "Can you speak Spanish, please?" we have a Directive act and the Request act which it abstracts. The hearer can again accept the Request and its implicatures. In the respective cases it can be proven that a Speech-Act occurred and that a Directive occurred, using the definition of abstraction.

An example in which ambiguity interferes with the planning process is this one. Suppose you are on a road trip with a friend, and as the friend is driving you pass a restaurant, and the friend says,

(76) Food?

The friend may be offering to stop at the restaurant, thinking that you may be

hungry. Or the friend may be hungry, and suggesting this possibility for food. You wish to reply promptly, in a way that accomodates each person's needs. A small amount of reasoning about who last ate what, in the appropriate belief space, may reveal to you that your friend is probably hungry even if you are not, so that you agree to stop. Or it may reveal that the friend expects just you to be hungry, and so you answer based on your own needs. This analysis is deeper than a simple implicature check in that there are many possible motivations for a suggestion, which an implicature check would not distinguish. Extended reasoning further restricts the range of interpretations.

Plan reasoning in this sense need not always be performed to reduce ambiguity. Vagueness and genuine ambiguity of intentions are quite common in speech and often not a problem. For instance, the speaker may mention plans to go to the store, and leave unclear whether this constitutes a promise.

In cases of genuine ambiguity, it is possible for the hearer to respond to each of the proposed interpretations, and indeed, politeness may even require it. Utterance (a) below could be a yes/no question *and* request, in a neutral context. Consider (b)-(j) as responses to (a). (We use ? and \* to indicate pragmatic appropriateness rather than grammaticality.)

- (77) a: Do you have our grades yet?  
 b: No, not yet.  
 c: No, I'm still working on them.  
 d: ?No, sorry.  
 e: ?No, I don't.  
 f: \*No.  
 g: Yes, I'm going to announce them in class.

- h: Sure, here's your paper. (hands paper.)
- i: Here you go. (hands paper.)
- j: \*Yes.

The most polite answers acknowledge the student's goal of knowing the grade; the least polite are the bare yes/no answers. This is not simply a question of shortness, because (d) and (e) are as long as (b). They don't provide any progress toward the student's goal. [Gibbs 86] claims that the very conventionality of "indirect" requests is related to their addressing the most likely obstacle to the request, but it is hard to be sure what this means. We claim that in cases like this, insisting on a single final labelling of the speech act is a mistake. The power of an indirect request lies precisely in its balancing of the two interpretations: if you won't satisfy the request you can answer the question in the negative, but what the asker is leading to is still recognizable.

Planning may be able to address multiple interpretations cheaply. The professor can easily plan responses which address both the goal of knowing whether the grades are ready and knowing what one's grade is.

- (78) a: No, but you did well as always.
- b: No, but it seems everyone missed number 6.
- c: Yes, I'm going to announce them in class.
- d: Sure, here's your paper.

In fact such examples are common. To "Can you speak Spanish?" one may reply "Si, si". To "Do you know what time it is?", giving the time implies that you know it. Or the hearer may simply ask what was meant.

Extended reasoning is thus invoked just when it may be of some help. One possible refinement would be to incorporate reasoning about misconceptions, such

as the work done by Pollack [Pollack 86]. Another refinement which will be helpful for real applications is to give the interpretations a weight, based on the implicature checking. One might, for instance, give every interpretation +1 for each true implicature and -2 for each false implicature. Such a scheme could make some misconceptions visible, because it allows interpretations to remain which have a false implicature but many other supporting ones. It should also be more robust. A sketch for such a scheme appears below.

give each interpretation +1 for true implicatures,

-2 for false.

now define some metrics to identify the different cases.

Then if there are none left with (say) positive ratings, invoke reasoning on the discards.

invoke reasoning on the discards.

If there is one left, or one significantly (say, by 2) favored over the others, take it and assert implicatures.

take it and assert implicatures.

If there is more than one with a positive rating, and their ratings are not significantly different, accept the ambiguity (&implicate?)

not significantly different, accept the ambiguity (&implicate?)

unless planning needs to resolve it, in which case use extended reasoning on the disjunction

on the disjunction

Further work is needed to make such a scheme practical.

### 6.3.2. An Example of Extended Reasoning

To see in detail how extended reasoning now works out novel interpretations, let us reconsider the example, "It's cold in here". The linguistic module generates these interpretations:

(INFORM AGENT A  
HEARER S  
PROP (COLD AREA Space1))

(SPEECH-ACT AGENT A)  
 HEARER S))

The plan reasoning is thereby constrained very little, since the entire range of speech acts is given. The next step is implicature checking. The context is this. Suppose you are in a car, by the only open window, and another passenger says "It's cold in here." Assume it's well known that a cold car causes the agent to be cold, that it is bad for agents to be cold, and that an open window can make the car cold. Further, assume that it is already well known that it is cold in the car.

MB(S, A, Attend(S, A))  
 MB(S, A, Attend(A, S))  
 ...  
 MB(S, A, Cold(space1))

The implicatures for the Inform act are

**Inform(A, S, Cold(Space1))**  
**Preconditions:** Attend(S, A),  
                   Attend(A, S)  
**Constraints:** B(A, Cold(Space1))  
                   Able(A, Self:Action),  
                   W(A, Self:Action)  
**Effects:** B(S, Cold(Space1))

Its implicatures are:

**preconditions hold:**  
 B(S, B(A, Attend(S, A)))  
 B(S, B(A, Attend(A, S)))

**constraints hold:**  
 B(S, B(A, B(A, Cold(Space1))))  
 B(S, B(A, Able(A, Inform(...))))  
 B(S, B(A, W(A, Inform(...))))

**effects do not hold:**  
 B(S, B(A, ~B(S, Cold(Space1))))

**effects intended:**



B(S, W(A, B(S, Cold(Space1))))

The problem with this interpretation is that its effects already hold. The hearer already believes that it is cold. Thus the interpretation is eliminated. The more abstract interpretation is

**Speech-Act(A, S)**  
**Preconditions:** Attend(S, A),  
                   Attend(A, S)  
**Constraints:** Able(A, Self:Action),  
                   W(A, Self:Action)

Its implicatures are:

preconditions hold:  
 B(S, B(A, Attend(S, A)))  
 B(S, B(A, Attend(A, S)))

constraints hold:  
 B(S, B(A, Able(A, Speech-Act(...))))  
 B(S, B(A, W(A, Speech-Act(...))))

These implicatures are all consistent, and the abstract interpretation remains. However, a single abstract interpretation is almost no interpretation at all, and so extended plan reasoning is invoked. The successful chain begins with the rejected Inform interpretation:

SBAW(INFORM(A, S, Cold(space1)))

SBAW(MB(S, A, AW(S KNOW Cold(space1)))) (action-effect)

SBAW(MB(S, A, AW(S KNOW Cold(A)))) (causal)

SBAW(MB(S, A, AW(S W not(Cold(A)))) (undesireability)

SBAW(MB(S, A, AW(S W not(Open(window1)))) (planning by causal)

SBAW(MB(S, A, AW(S W Close(S,window1)))) (planning by effect-action)

SBAW(MB(S, A, AW(Close(S,window1))) (want-action)

SBAW(Request (A, S, Close(S>window1))) (body-action)

The only difference between this chain of reasoning and the original example (see Chapter Four) is that it begins with the rejected Inform interpretation rather than an S-Inform. After it is completed, the implicatures for the interpretation must be calculated, checked, and asserted along with the interpretation. In an example with more restricted output from the linguistic module, we would see that many possible chains of reasoning are eliminated when they arrive at an interpretation outside of that interpretation range.

## 7. Implementation

This chapter describes the computer programs which embody the ideas in this dissertation. Two main components of this dissertation have been implemented. The first is the incremental, linguistic component, which we will refer to as the linguistic module. It is written in Common Lisp and runs on Symbolics LISP machines, on UNIX workstations. The second computes the implicatures of the suggested speech act interpretations, and we will refer to this component as the implicature module. It makes use of the Rhetorical (Rhet) knowledge representation system, and therefore is only available on the Symbolics. Together, these two modules effectively handle examples like the Spanish example, which requires no long-chain reasoning. Extended plan reasoning is an unimplemented third module. These modules can be loaded as components of the Rochester Discourse System [Allen 89], or individually.

### 7.1. The Rochester Discourse System

The discourse system is a study in the integration of several aspects of discourse processing. These aspects may be well understood individually, but no previous system has successfully combined them. This system achieves integration of its components via a so-called blackboard architecture [Lesser 77]. A blackboard is a global data structure which all modules can read and to which they all may write, allowing them to proceed asynchronously of each other. Each module definition includes patterns specifying its desired input, and the module is invoked (subject to

some scheduling algorithm) when matching input appears on the blackboard. It writes its output to the blackboard as well.

The blackboard is divided into parts. One portion is the chart used by the chart parser and semantic interpretation. Detailed partial interpretations appear here which do not concern other modules. The remainder is divided into *segments*, intersentential units of discourse structure [Grosz 86a]. A possible speech act interpretation may be tentatively associated with several different segments, and within them, with each of several choices for any other structure, such as referents of a noun phrase. This web of possibilities is developed using best-first search. Our linguistic and implicature modules are invoked essentially from within a segment, on a sentence interpretation which may contain disjunctions. As yet they say nothing about what speech acts may continue a segment, nor do they rate possible interpretations to guide search.

The discourse system's nonlinguistic world knowledge is managed by the Rhetorical knowledge representation system [Miller 87]. Rhet is a Horn-clause theorem prover supporting not only forward and backward chaining, but also advanced features such as structured types, reasoning about equality, various proof modes, the Tempos time reasoner [Koomen 88, Koomen 89], and a hierarchy of belief spaces. The discourse system is loadable in Rhet, making Rhet's program interface available to the modules with no restrictions. The implicature module shares the knowledge base with reference, plan reasoning and other high-level modules, and all of

these modules make demands on the plan hierarchy stored there. Neither Rhet nor the blackboard understand each other's data structures, although they can store them. This means that modules using both must perform translations between the two.

## 7.2. The Linguistic Module

This module generates speech act interpretations of utterances, by the method described in chapters 2 and 3. It is similar to a bottom-up parser, taking as input a sentence representation containing lexical, syntactic, logical form, and reference information. It matches a set of patterns against this utterance, combining the results incrementally. It then generates action descriptions for the utterance, interpreting it as a set of possible speech acts. It requires for operation both a set of interpretation rules and a set of action definitions.

For simplicity, the implementation assumes that the utterance has already undergone semantic interpretation and reference analysis. This ensures that all information that might be needed to construct the speech act interpretation is already available. (This should be reworked for experiments in control flow.) The pattern on which the blackboard invokes the module is thus very simple; it specifies any utterance or part, with syntactic unit S (clause), which has already been assigned some logical form and knowledge base structure. Its declaration is shown below.

```
(BBDeclarePattern 'GENERATOR
  '(ENTRY %id (SYN (S %s) (SEM %sem) (REF %ref)))
  'SSA)
```

In theory, many interpretations can be done with much less information, and on much weaker syntactic structures. When it is invoked by the blackboard, the module translates the discourse system's representation of the structure into its own internal representation. This representation can then be matched against the speech act interpretation rules.

The module's internal representation of utterances is a LISP list, consisting of a category symbol followed by any number of slot/filler pairs. This is very much as described in chapter 2. A category may be a syntactic category or feature like NP, a logical form class like Capable, or a knowledge base type like Inform. A slot/filler pair is also a list, consisting of a slotname followed by a word or some other value, or a (category (slot filler) ...) structure. Names of Rhet objects have square brackets. Here is the sentence "Can you go to the store?":

```
(setq s1 '(s (mood y-n-q)
  (voice act)
  (subj (np (pro you)
    (sem h1)
    (ref [H]))))
  (auxs can)
  (main-v go)
  (tense pres)
  (mods (pp (prep to)
    (pobj (np (det the)
      (head store)
      (sem (STORE (id sto1)
        (num 1)
        (gen n)))
      (ref [store7])))))
  (sem (CAPABLE (time pres)
    (agent h1)
    (theme (GO (to-loc sto1))))))
  (ref [able1])))
```

The linguistic module is based on a simple bottom-up parser. The parser takes as input a sentence representation, as shown above, and a set of rules. Here is a small set of rules:

```
(( $ ) (T-SPEECHACT (R-AGENT [S] )))
(( $ ) (T-SPEECHACT (R-HEARER [H] )))

; please signals a request
(( $ (ADV PLEASE))
  (T-REQUEST (R-OBJECT (V OBJ REF)) )
)

; "can you....?" may be a request or some other act
(( $ (AUXS /MODALS) (MOOD Y-N-Q)
  (VOICE ACT) (SUBJ (NP (PRO YOU)))
)
  (T-REQUEST (R-OBJECT (V OBJ REF)))
  (T-SPEECHACT)
)

; a yes-no question may be a yes-no question or some other act
(( $ (MOOD Y-N-Q)
  (T-ASK (R-PROP (V REF)))
  (T-SPEECHACT)
)
)

(setq /MODALS '(CAN COULD WOULD WILL MIGHT))
```

The rules use the representation discussed above, with a few extensions. They are lists, containing a left hand side followed by all of the possible interpretations. The left hand side of the first two rules consists simply of \$, a wildcard matching any category. They thus match any syntactic unit, and simply fill in the speaker and hearer. [S] and [H] should already be bound in context using the BB functions retrieving speaker and hearer.

The third rule says that any syntactic unit to which the adverb "please" is attached is a request. The requested object is found in the object role of the sentence's knowledge base interpretation, using the value function. All lists beginning with V are interpreted as the value function: the remainder of the list is a series of slot names, which the value function uses to retrieve information from the utterance. It simply uses the slot names to trace down into the utterance's structure, and returns the contents of the deepest one, filling out the new structure being created.

The fourth and fifth patterns are more complicated, and have more than one speech act interpretation. For a given domain, such rules could be weighted according to their predictive power, and the output given corresponding weights on the blackboard. Atoms beginning with / are names of disjunctive lists, so rule four matches sentences containing any of the five modal auxiliaries listed in /MODALS.

For a rule to match a structure, the category of the left hand side must be the category of the structure. Each of the slots in the left hand side must be present in structure too, with the same value. The structure may have extra, unmatched slots. We have already discussed the role of wildcards and disjunction. The utterance structure given above matches the wildcard in rules one and two, yielding the first two interpretations below. It does not have an ADV slot containing "please", so it fails to match the third rule. It has the correct category and mood for the fourth rule, one of the list of modal verbs, and a subject that matches recursively. This match yields two possible interpretations, with the object of the request the object of [able1], [go881]. The final match is a simple match on sentence mood.



```

((T-SPEECHACT (R-AGENT [S])))
((T-SPEECHACT (R-HEARER [H])))
((T-REQUEST (R-OBJECT [go881]))
 (T-SPEECHACT))
((T-ASK [able1])
 (T-SPEECHACT))

```

Those are the four sets of top-level matches for our example sentence and rule set. However, the pattern-matcher is based on a bottom-up parser and therefore performs this whole process at each level. The parser descends the sentence representation recursively, then backs out, attempting to apply the rule set at each level. If a rule matches, the corresponding interpretation is generated. This partial interpretation is merged with any others before backing out to the next level, and the corresponding action instance in the knowledge base is created. The merge operation is to take the cross product of the sets, and combine the interpretations in the resulting sets. The combining is unification or graph matching, in which categories intersect, slots union, and slot values intersect. For the rules and utterance given above, these are the interpretations:

```

(T-REQUEST (R-AGENT [S])
            (R-HEARER [H])
            (R-OBJECT [go881]))
(T-ASK (R-PROP [able1])
        (R-AGENT [S])
        (R-HEARER [H]))
(T-SPEECHACT (R-AGENT [S])
              (R-HEARER [H]))

```

The set containing both Request and Ask interpretations cannot undergo combination, and is eliminated. The set of three interpretations is returned to the blackboard in a new slot on the utterance:

```
(BBDefineSlotValue utterance 'SAs value-list context)
```

Alternatively, the implicature module can be invoked directly on the list.

### 7.3. Implicature

The implicature module takes as input a list of possible speech act interpretations as generated by the linguistic module. The blackboard invokes it on this pattern definition:

```
(BBDeclarePattern 'GENERATOR
                  '(ENTRY %id (SYN (SAs %s)))
                  'IMPL)
```

This insists only that the list of interpretations be precomputed by the linguistic module. They must, however, be translated into the Rhetorical knowledge representation before their implicatures can be computed.

The Rhetorical knowledge representation language is first order predicate calculus with typed objects. It has been extended to include a framelike language of structured objects which have roles that can be filled by objects and are inherited via an abstraction hierarchy. Propositions can be associated with types Initializations, Constraints, or Relations; the first two processed by the system and the third for interpretation by user programs. We use all of these facilities to

represent plans in Rhet.

The hierarchy is rooted at the universal type, T-U. Four kinds of Relations are defined for plans: preconditions, constraints, effects, and body.

get a better example

```
(Define-Subtype 'T-Language 'T-U)
```

```
(Define-Instance [English] 'T-Language)
```

```
(Define-Subtype 'T-SpeechAct 'T-Plan
```

```
  :Roles '((R-hearer T-Human) (R-Language T-Language))
```

```
  :Initializations '([Set-Function-Value [F-Language ?self] [English]]
```

```
    [Set-Function-Value [F-Preconditions ?self]
```

```
      ([Listening [F-Hearer ?self]]
```

```
        [Noise-Free-Line]))
```

```
    [Set-Function-Value [F-CONSTRAINTS ?self]
```

```
      ([Speaks [F-Agent ?self] [F-Language ?self]]
```

```
        [Speaks [F-Hearer ?self] [F-Language ?self]])))]
```

```
(Define-Functional-Subtype 'T-Request 'T-SpeechAct
```

```
  :Roles '((R-object T-Plan))
```

```
  :Initializations '([Set-Function-Value [F-EFFECTS ?self]
```

```
    ([Exec [F-Hearer ?self] [F-Object ?self]] ) )
```

```
    [Set-Function-Value [F-Constraints ?self]
```

```
      ([Able [F-Hearer ?self] [F-Object ?self]] ) ) )
```

A SpeechAct has the inherited role Agent as well as two of its own. It requires the speaker and hearer to share a specific language, and for the hearer to be listening. A Request more specifically has a requested object, which is an action. A constraint is that the hearer is able to do the action, and the effect is that they actually do.

One major caveat of this representation is that it cannot express second order constructs like the type Inform(S, H, P) where P is a proposition. We dodge this by creating a structured type for the head of P, so that P becomes a function term.

Then, in order to use Rhet as the database for these predicate-objects, we invent a TRUE predicate and a DO predicate, and so on, which apply to the predicate-objects and can be proved by Rhet. Since we do not provide a full second-order language with negation &c, we cannot move negation, belief, and other modal operators and functions (?and, or) over the TRUE predicate.

The implicature computation algorithm makes heavy use of the context mechanisms provided by Rhet. There are two of these: belief contexts and user contexts. Belief contexts form a tree whose root is the database of all things mutually believed by all agents being modelled, and whose leaves are the databases modelling beliefs particular to one agent. Intermediate databases such as SBHBMB are created as necessary. And most implicatures are in the SBHB category to start with. A leaf database inherits the contents of all databases on its path to the root. User contexts may be created which inherit from any chosen belief context.

The implicature algorithm takes the list of speech act interpretations and for each, creates a user context beneath SBHB. It translates each interpretation into the knowledge representation, creating an action instance of the appropriate speech act type, with variable bindings and in the corresponding context. Thus for each interpretation, retraction is cheap and further reasoning can be done. The further reasoning consists of the consistency checking via the implicatures: A procedure takes each predicate, say, each precondition, and builds and checks each corresponding implicature. The check itself is not a full superexponential proof of

database consistency, but rather an attempt to prove both the implicature and its negation. If neither succeeds, the implicature is asserted in this context. If all implicatures are proveable or asserted, the speech act interpretation is returned along with its context. Otherwise the interpretation is inconsistent and its context is destroyed. (For correcting misconceptions this context should be saved for further analysis.) The output of the algorithm is thus a list of speech acts and their contexts, from which interpretations implicating certain obvious contradictions have been eliminated.

The actual implicatures computed are the ones based on preconditions, constraints, and negations of effects. The effects' being intended was not implemented because although there is a hierarchy of belief spaces in Rhet, there is no corresponding set of intention spaces.

#### 7.4. Limitations

The linguistic module constructs speech act interpretations incrementally by matching its rules against an input structure. The implicature module filters such sets based on three of the four classes of implicatures we have discussed.

There are of course open problems. One would like to experiment with large interpretation rule sets, and with the constraints from other modules. In addition the RHETORICAL plan representation has changed significantly, so that it would be desirable to reimplement the implicature component from scratch.

## 8. Conclusion

As a measure of our progress, let us reconsider the first example of this document.

A is standing by an obviously immobilized car and is approached by B;  
the following exchange takes place:

- (1) A: I am out of petrol.  
B: There is a garage round the corner.

B communicates that the garage is open and has gas to  
sell, and so on.

We now have a mechanism whereby hearers recognize each other's intentions, using both linguistic and general reasoning ability. A's utterance is a request for help, if A so intends it and if B is able to use our mechanism to identify this intention. In this particular case, B might identify a request either via the incremental, semantic "I need" rule, or by inference from the context. Its speech-act based implicatures are plausible in this context. B's helpful suggestion may likewise be so identified by A. Note, however, that the speech-act based implicatures which we have used for screening are not precisely the ones listed by Grice for this example. Those listed by Grice do indeed bear a strong resemblance to preconditions, constraints, and so on, but they are preconditions of the domain plan to buy gas. Such conclusions are indeed cancellable and detachable, as discussed in [Hinkelman 87]. They require knowledge of the agents' goals and plans for their interpretation. They are also plan-based conversational implicatures.

Plan-based conversational implicatures merit further study. In particular, one would like to see how cancellation mechanisms operate on them, and how a speech

act interpretation can possibly be accepted with cancelled implicatures. Would the speech-act computation be affected, or would the cancellation mechanism be able to operate successfully just on the results? As it is, the implicatures partially determine the speech act interpretation. If the interpretation algorithm is altered further to accommodate the cancellation process, speech acts and implicatures will be tightly bound indeed. The original concepts of speech act and implicature were never declared to be disjoint; here we have made one suggestion about where to make a cut.

It should be noted that our implicature calculations do not make use of explicit representations of Gricean maxims. Rather, the mechanisms simply operate in accordance with Gricean principles. There are many further questions about our mechanism to answer.

Are speech act classes idioms? In Chapter One we argued that they are not merely fixed lexical strings, nor are they merely rigid semantic structures. Rather, conventional speech acts are linguistic structures matching a pattern or constellation of linguistic features, and which have certain pragmatic consequences. If we wish to regard conventional acts as (pragmatic) idioms in this richer sense, as Sadock does [Sadock 74], we are led immediately to the suggestion that idioms themselves are patterns of lexical, syntactic, and semantic features with certain semantic consequences. Idioms under such a theory would be more than oversized lexical entries; they would have significant internal structure, and implications for the architecture of natural language processors. While perhaps not internally

maintaining compositionality of meaning, they would participate in meaning incrementally and make allowances for the great permeability of some idioms to substitutions: "As X as a Y", for example. As does our theory of speech acts, such a theory of idioms would imply that NL architectures must either permit semantic and pragmatic processes very early access to lexical input, or make that input available later at the point where those processes are invoked. In sum, conventional speech acts are idioms in an interesting sense of the word.

Robustness and scalability are issues common to all rule-based systems. Both the linguistic and plan inference components can accommodate addition of new rules chosen from a very general class. Together, they handle a wider range of phenomena than previously. It remains to be shown that rule sets large enough for detailed linguistic coverage are still reliable and modifiable. (A large rule set might be a hundred or two rules, compared with thousands for many working grammars.) For this an attempt should be made to handle some extended corpus of dialogue. The implicature rules are designed to localize search, but this will be effective in large databases only if they can be indexed appropriately. The range of speech act types discussed has been fairly broad, and limited primarily by knowledge representation issues. Progress in the representation of intention and other basic concepts, and representation of physical and social activities, will greatly improve the speech act definitions. Representation of speech act definitions is also complicated by the addition of acts with discourse control functions (see also [Mann 88]). Since representation of discourse structure is an active area of



research, it may soon be possible to handle these kinds of speech acts too. Thus we expect our system to be more robust than previous proposals, but would like to verify that it scales up.

We have not solved the problem of control associated with extended reasoning. We avoid it as much as possible by emphasizing the locality of the implicature checking process. An underlying issue here is resource allocation, which is gaining some credibility in the literature of knowledge representation [Perlis 89] .

We have essentially assumed that traditional lexical, syntactic, and semantic analysis were practical, including reference. Many linguistic issues remain, however. The problem of reference is far from solved, and plan-based speech acts provide a set of constraints which may be helpful in identifying referents of expressions. The problem would then be to allow these two processes to interact in a satisfactory manner. Another linguistic problem is that certain constructs in English have especial influence on speech act interpretation: mood, modal verbs, adverbs, and adverbial phrases. We have treated them simply as various resources which can signal particular intentions, since they show many irregularities. Yet they may have a few more generalizations to offer, especially with regard to speech act interpretation. Studies of conjoined speech acts in English, and cross-linguistic studies of speech acts, are obvious next steps.

In summary, to determine what an agent is doing by making an utterance, we must make use of not only general reasoning about actions in context, but also the

linguistic features which by convention are associated with specific speech act types. To do this, we match patterns of linguistic features as part of the standard linguistic processing. The resulting partial interpretations are merged, and then filtered by determining the plausibility of their conversational implicatures. If there is not a unique plausible interpretation, full plan reasoning is called. Remaining ambiguity is not a problem but simply a more complex basis for the hearer's planning processes. Linguistic patterns and plan reasoning together constrain speech act interpretation sufficiently for discourse purposes.

## References

- [Allen 83] Allen, J., "Recognizing Intentions From Natural Language Utterances," in *Computational Models of Discourse*, Brady, M. and Berwick, B. (ed.), MIT Press, Cambridge, MA, 1983, 107-166.
- [Allen 87] Allen, J., *Natural Language Understanding*, Benjamin Cummings Publishing Co., 1987.
- [Allen 89] Allen, J. F., Guez, S., Hoebel, L., Hinkelman, E., Kyburg, A., Jackson, K. J. and Traum, D., "The Discourse System Project," Technical Report 317, Department of Computer Science, University of Rochester, Rochester, NY, 1989.
- [Appelt 85] Appelt, D., *Planning English Sentences*, Cambridge University Press, 1985.
- [Austin 62] Austin, J. L., *How to Do Things with Words*, Harvard University Press, Cambridge, MA, 1962.
- [Brown 80] Brown, G. P., "Characterizing Indirect Speech Acts," *American Journal of Computational Linguistics* 6:3-4, July-December 1980, 150-166.
- [Carberry 87] Carberry, S., "Plan Recognition in User Modelling," *Computational Linguistics*, 1987.
- [Carlson 84] Carlson, G. N., "Thematic Roles and their Role in Semantic Interpretation," *Linguistics* 22, 1984, 259-279.
- [Cohen 86] Cohen, P. R. and Levesque, H. J., "Persistence, Intention, and Commitment," *Proc. 1986 Timberline Workshop on Planning and Practical Reasoning*, Los Altos, 1986.
- [Cohen 88] Cohen, P. R. and Levesque, H. J., *Rational Interaction as the Basis for Communication*, unpublished manuscript, October 1988.
- [Cullingford 85] Cullingford, R., "SAM," in *Natural Language Processing*, Grosz, B. J., Jones, K. S. and Webber, B. L. (ed.), Morgan Kaufmann, Los Altos, CA, 1986, 627-650.
- [Davison 83] Davison, A., "Linguistic or Pragmatic Description in the Context of the Performatadox," *Linguistics and Philosophy* 6, 1983, 499-526, D. Reidel Publishing Company.
- [Fikes 71] Fikes, R. E. and Nilsson, N. J., "STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving," *Artificial Intelligence* 2:3/4, 1971, 189-208.
- [Gibbs 84] Gibbs, R., "Literal Meaning and Psychological Theory," *Cognitive Science* 8, 1984, 275-304.
- [Gibbs 86] Gibbs, R. W., "What Makes Some Indirect Speech Acts Conventional?," *Journal of Memory and Language* 15, 1986, 181-196, Academic Press, Inc..
- [Goldman 70] Goldman, A. I., *A Theory of Human Action*, Prentice-Hall, Englewood Cliffs, NJ, 1970.

- [Gordon 75] Gordon, D. and Lakoff, G., "Conversational Postulates," in *Syntax and Semantics V. 3*, Cole, P. and Morgan, J. L. (ed.), Academic Press, New York, 1975.
- [Grice 57] Grice, H. P., "Meaning," *Philosophical Review* 66, 1957, 377-388.
- [Grice 71] Grice, H. P., "Meaning," in *Semantics*, Steinberg, D. D. and Jakobovits, L. A. (ed.), Cambridge University Press, Cambridge, 1971, 53-60.
- [Grice 75] Grice, H. P., "Logic and Conversation," in *Syntax and Semantics: Speech Acts*, Cole, P. and Morgan, J. L. (ed.), Academic Press, New York, 1975, 41-58.
- [Grosz 86a] Grosz, B. J. and Sidner, C. L., "Attention, Intentions, and the Structure of Discourse," *Computational Linguistics* 12:3, July-September 1986, 175-204.
- [Grosz 86b] Grosz, B. J., "The Representation and Use of Focus in a System for Understanding Dialogues," in *Natural Language Processing*, Grosz, B. J., Jones, K. S. and Webber, B. L. (ed.), Morgan Kaufmann, Los Altos, CA, 1986, 353-362.
- [Grosz 87] Grosz, B. J. and Sidner, C. L., *Plans for Discourse*, Symposium on Intentions and Plans in Communication and Discourse, 1987.
- [Hinkelman 87] Hinkelman, E., "Thesis Proposal: A Plan-Based Approach to Conversational Implicature," Technical Report 202, Dept. Computer Science, University of Rochester, June 1987.
- [Hirschberg 85] Hirschberg, J., "A Theory of Scalar of Implicature," MS-CIS-85-56, PhD Thesis, Department of Computer and Information Science, University of Pennsylvania, December 1985.
- [Horn 84] Horn, L. R. and Bayer, S., Short-Circuited Implicature: A Negative Contribution, Vol. 7, 1984.
- [Horrigan 77] Horrigan, M. K., "Modelling Simple Dialogs," Technical Report 108, Master's Thesis, Department of Computer Science, University of Toronto, Toronto, May 1977.
- [Jackendoff 72] Jackendoff, R. S., *Semantic Interpretation in Generative Grammar*, MIT Press, Cambridge, 1972.
- [Kautz 87] Kautz, H. A., "A Formal Theory of Plan Recognition," TR215, Department of Computer Science, University of Rochester, Rochester, NY, May 1987.
- [Koomen 88] Koomen, J. A. G. M., *The TIMELOGIC Temporal Reasoning System*, Department of Computer Science, University of Rochester, Rochester, NY, October 1988. TR231.
- [Koomen 89] Koomen, J. A. G. M., *Reasoning About Recurrence*, Department of Computer Science, University of Rochester, Rochester, NY, August 1989.
- [Leech 75] Leech, G. and Svartvik, J., *A Communicative Grammar of English*, Longman, Harlow, 1975.
- [Lehnert 78] Lehnert, W., *The Process of Question Answering*, Lawrence Erlbaum Associates, Hillsdale, NJ, 1978.

[Lesser 77] Lesser, V. R. and Erman, C. D., "A Retrospective View of Hearsay-II Architecture," *Proc. Fifth International Joint Conference on Artificial Intelligence*, August 1977, 790-800.

[Lightfoot 79] Lightfoot, D. W., *Principles of Diachronic Syntax*, Cambridge University Press, Cambridge, 1979.

[Litman 85] Litman, D. J., "Plan Recognition and Discourse Analysis: Integrated Approach for Understanding Dialogues," Technical Report 1, Department of Computer Science, University of Rochester, Rochester, NY, 1985.

[Mann 88] Mann, W. C., "Dialogue Games: Conventions of Human Interaction," *Argumentation* 2, 1988, 511-532, Kluwer Academic Publishers.

[McCafferty 86] McCafferty, A. S., Explaining Implicatures, 23 October 1986.

[McKeown 86] McKeown, K. R., "Discourse Strategies for Generating Natural Language Text," in *Natural Language Processing*, Grosz, B. J., Jones, K. S. and Webber, B. L. (ed.), Morgan Kaufmann, Los Altos, CA, 1986, 479-500.

[Miller 87] Miller, B. and Allen, J., *The Rhetorical Knowledge Representation System: A User's Manual*, forthcoming technical report, Department of Computer Science, University of Rochester, 1987.

[Morgan 75] Morgan, J. L., "Some Interactions of Syntax and Pragmatics," *Syntax and Semantics 3: Speech Acts*, 1975, 289-303, Academic Press.

[Nilsson 80] Nilsson, N. J., *Principles of Artificial Intelligence*, Tioga Publishing Company, Palo Alto, CA, 1980.

[Pelavin 86] Pelavin, R. and Allen, J. F., "A Formal Logic of Plans in Temporal Rich Domains," *Proceedings of the IEEE* 74:10, October 1986, 1364-1382.

[Perlis 89] Perlis, D., "Thing and Thought," *Proc. Society for Exact Philosophy*, 1989.

[Perrault 78] Perrault, C. R., Allen, J. F. and Cohen, P. R., "Speech Acts as Basis for Understanding Dialogue Coherence," in *Proc. of Second Conference on Theoretical Issues in Natural Language Processing*, Champaign-Urbana, 1978.

[Perrault 80] Perrault, C. R. and Allen, J. F., "A Plan-Based Analysis of Indirect Speech Acts," *American Journal of Computational Linguistics* 6:3-4, July-December 1980, 167-82.

[Perrault 87] Perrault, C. R., "An Application of Default Logic to Speech Act Theory," in *SDF Benchmark Series: Plans and Intentions in Communication and Discourse*, Cohen, P. R., Morgan, J. and Pollack, M. (ed.), MIT Press, Cambridge to appear 1987.

[Pollack 86] Pollack, M. E., "Inferring Domain Plans in Question-Answering," MS-CIS-86-40, PhD Thesis, Department of Computer and Information Science, University of Pennsylvania, May 1986.

[Reiter 78] Reiter, R., "On Reasoning by Default," *TINLAP-2*, New York, NY 1978, 210-218.

[Sacerdoti 74] Sacerdoti, E. D., "Planning in a Hierarchy of Abstraction Spaces," *Artificial Intelligence* 5, 1974, 115-135.

- [Sacerdoti 80] Sacerdoti, E. D., "Plan Generation and Execution for Robotics," Technical Note 209, SRI International, Menlo Park, CA, April 1980.
- [Sadock 74] Sadock, J. M., in *Toward a Linguistic Theory of Speech Acts*, Academic Press, New York, 1974.
- [Sadock ng] Sadock, J. M., Comments on the Papers by Vanderveken and Cohen and Levesque, forthcoming.
- [Schank 77] Schank, R. C. and Abelson, R., *Scripts, Plans, Goals, and Understanding*, Lawrence Erlbaum Associates, Hillsdale, NJ, 1977.
- [Searle 75] Searle, J., "Indirect Speech Acts," in *Syntax and Semantics, v3: Speech Acts*, Cole and Morgan (ed.), Academic Press, New York, NY, 1975.
- [Searle 79] Searle, J. R., "A Taxonomy of Illocutionary Acts," in *Expression and Meaning*, Cambridge University Press, London, 1979, 1-29.
- [Sidner 81] Sidner, C. L. and Israel, D. J., "Recognizing Intended Meaning and Speakers' Plans," *Proc. IJCAI '81*, 1981, 203-208.
- [Tenenberg 89] Tenenberg, J. D., "Extending Inheritance Abstraction to Symbolic Planning Systems," *Research Review*, Rochester, NY, 1989.
- [Wilensky 83] Wilensky, R., in *Planning and Understanding*, Addison-Wesley Publishing Company, Reading, MA, 1983.